

Parageometric outer automorphisms of free groups

Michael Handel* and Lee Mosher†

February 1, 2008

Abstract

We study those fully irreducible outer automorphisms ϕ of a finite rank free group F_r which are *parageometric*, meaning that the attracting fixed point of ϕ in the boundary of outer space is a geometric \mathbf{R} -tree with respect to the action of F_r , but ϕ itself is not a geometric outer automorphism in that it is not represented by a homeomorphism of a surface. Our main result shows that the expansion factor of ϕ is strictly larger than the expansion factor of ϕ^{-1} . As corollaries (proved independently by Guirardel), the inverse of a parageometric outer automorphism is neither geometric nor parageometric, and a fully irreducible outer automorphism ϕ is geometric if and only if its attracting and repelling fixed points in the boundary of outer space are geometric \mathbf{R} -trees.

1 Introduction

There is a growing dictionary of analogies between theorems about the mapping class group of a surface $\mathcal{MCG}(S)$ and theorems about the outer automorphism group of a free group $\text{Out}(F_r)$. For example, the Tits alternative for $\mathcal{MCG}(S)$ [McC85] is proved using Thurston's theory of measured geodesic laminations [FLP⁺79], and for $\text{Out}(F_r)$ it is proved using the Bestvina–Feighn–Handel theory of laminations ([BFH97], [BFH00], [BFH05]).

Expansion factors. Here is a result about $\mathcal{MCG}(S)$ of which one might hope to have an analogue in $\text{Out}(F_r)$. Given a finitely generated group G , its outer automorphism group $\text{Out}(G)$ acts on the set of conjugacy classes \mathcal{C} of G . Given $c \in \mathcal{C}$ let $\|c\|$ be the smallest word length of a representative of c . Given $\phi \in \text{Out}(G)$ define the *expansion factor*

$$\lambda(\phi) = \sup_{c \in \mathcal{C}} \left(\limsup_{n \rightarrow +\infty} \|\phi^n(c)\|^{1/n} \right)$$

*Supported in part by NSF grant DMS0103435.

†Supported in part by NSF grant DMS0103208.

If $\phi \in \mathcal{MCG}(S) \approx \text{Out}(\pi_1 S)$ is pseudo-Anosov then $\lambda(\phi)$ equals the pseudo-Anosov expansion factor [FLP⁺79]. By combining results of Thurston and Bers one obtains:

Theorem. *If $\phi \in \mathcal{MCG}(S)$ is pseudo-Anosov, then there is a unique ϕ -invariant geodesic in Teichmüller space, consisting of the points in Teichmüller space which minimize the translation distance under ϕ . This translation distance equals $\log(\lambda(\phi))$.*

Consider a free group F_r and the Culler-Vogtmann *outer space* \mathcal{X}_r on which $\text{Out}(F_r)$ acts properly. Recall that $\phi \in \text{Out}(F_r)$ is *reducible* if there is a nontrivial free decomposition $F_r = A_1 * \cdots * A_k * B$ such that ϕ permutes the conjugacy classes of A_1, \dots, A_k ; otherwise, ϕ is irreducible. If ϕ^k is irreducible for all $k \geq 1$ then we say that ϕ is *fully irreducible*. By analogy, a mapping class on S is reducible if it preserves the isotopy classes of some nontrivial decomposition of S into essential subsurfaces, and is fully irreducible if and only if it is represented by a pseudo-Anosov homeomorphism.

Question: Does the above theorem have an analogue for a fully irreducible $\phi \in \text{Out}(F_r)$? The question does not quite make sense because no metric is specified on \mathcal{X}_r , but one can instead ask: Is there an analogue with respect to *some* $\text{Out}(F_r)$ -equivariant metric on \mathcal{X}_r ? Or on any other metric space on which $\text{Out}(F_r)$ acts?

Answer: No. If translation distance for ϕ is uniquely minimized on an axis γ then, by symmetry of the distance function, translation distance for ϕ^{-1} is also uniquely minimized on γ , and the minima for ϕ and for ϕ^{-1} along γ are equal. By equating minimal translation distance with $\log(\lambda)$, one would conclude that $\lambda(\phi) = \lambda(\phi^{-1})$. However, an example from [BH92] has the property that $\lambda(\phi) \neq \lambda(\phi^{-1})$: consider $\phi \in \text{Out}(F_3)$ and ϕ^{-1} represented by the automorphisms

$$\Phi: \begin{cases} A & \rightarrow AC \\ B & \rightarrow A \\ C & \rightarrow B \end{cases} \quad \Phi^{-1}: \begin{cases} A & \rightarrow B \\ B & \rightarrow C \\ C & \rightarrow \overline{B}A \end{cases}$$

Interpreting these formulas as self maps of the three-petaled rose, each is clearly a train track map. From the results of [BH92] (see also Proposition 4), if $f: G \rightarrow G$ is a train track representative of $\phi \in \text{Out}(F_r)$, and if $\text{PF}(M_f)$ denotes the Perron-Frobenius eigenvalue of the transition matrix M_f , then $\lambda(\phi) = \text{PF}(M_f)$. For the above two train track maps we therefore obtain

$$\lambda(\phi) = \text{PF} \begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} > 1.4, \quad \lambda(\phi^{-1}) = \text{PF} \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} < 1.4$$

Confronted with such a strange phenomenon, one strategy is to see what appropriately weaker results can be proved. We follow this strategy in the companion

paper [HM06b] where we show that the ratio $\log(\lambda(\phi))/\log(\lambda(\phi^{-1}))$ is bounded by a constant depending only on the rank r . This is what one would expect if there were an axis γ for ϕ with translation distance $\log(\lambda(\phi))$ and an axis γ' for ϕ^{-1} with translation distance $\log(\lambda(\phi^{-1}))$, such that γ, γ' are fellow travelers. Encouraged by this result, we have pursued the study of axes in outer space, with some interesting analogues of uniqueness of axes [HM06a].

Parageometric outer automorphisms. In this work we pursue another strategy: explore the strange phenomenon on its own terms. The ϕ described above turns out to be an example of a *parageometric* outer automorphism, as we discovered by comparing discussions of this same example in [BF95] and in [BF]. The interest in this concept was pointed out in [GJLL98], where we found the terminology “parageometric”.

While we believe that the phenomenon $\lambda(\phi) \neq \lambda(\phi^{-1})$ is generic among fully irreducible outer automorphisms ϕ , we shall show that inequality *always* holds when ϕ is parageometric, in fact we give an explicit geometric argument which shows that $\lambda(\phi) > \lambda(\phi^{-1})$.

To define parageometricity, recall the action of $\text{Out}(F_r)$ on the compactified outer space $\overline{\mathcal{X}}_r = \mathcal{X}_r \cup \partial\mathcal{X}_r$ consisting of (classes of) very small actions of F_r on \mathbf{R} -trees (see [CL95] for simplicial \mathbf{R} -trees and [BF] for nonsimplicial \mathbf{R} -trees); we shall call these objects “ F_r -trees”. The action of a fully irreducible $\phi \in \text{Out}(F_r)$ on $\overline{\mathcal{X}}_r$ has source–sink dynamics, with a repelling F_r -tree $T_- \in \partial\mathcal{X}_r$, and an attracting F_r -tree $T_+ \in \partial\mathcal{X}_r$ (see [BFH97] for orbits in \mathcal{X}_r and [LL03] for orbits in $\partial\mathcal{X}_r$). An F_r -tree is *geometric* if it is dual in the appropriate sense to a measured foliation defined on some 2-complex whose fundamental group surjects to F_r [LP97]. For example, if a fully irreducible $\phi \in \text{Out}(F_r)$ is geometric, meaning that it is represented by an automorphism of a surface with boundary, then both of the F_r -trees T_-, T_+ are geometric: this follows from Thurston’s theorem that ϕ is represented by a pseudo-Anosov surface homeomorphism $f: S \rightarrow S$, because T_- and T_+ are dual to the stable and unstable measured foliations of f defined on the surface S .

A fully irreducible $\phi \in \text{Out}(F_r)$ is said to be *parageometric* if T_+ is a geometric F_r -tree but ϕ is not a geometric outer automorphism.

For example, the outer automorphism ϕ described above is parageometric: geometricity of the F_r -tree T_+ is proved in Example 3.4 of [BF]; and Levitt’s “thinness” property for T_- is proved in Example 10.1 of [BF95], showing that T_- is not a geometric F_r -tree, and so ϕ is not a geometric outer automorphism. In Proposition 6 we will gather results of [BF] and [BH92] which give a method of characterizing parageometricity solely from the properties of a train track representative.

Here are our main results:

Theorem 1. *If $\phi \in \text{Out}(F_r)$ is parageometric then $\lambda(\phi) > \lambda(\phi^{-1})$.*

Corollary 2. *If $\phi \in \text{Out}(F_r)$ is parageometric then ϕ^{-1} is neither geometric nor parageometric.*

Proof. If ϕ^{-1} is geometric then $\lambda(\phi^{-1}) = \lambda(\phi)$, whereas if ϕ^{-1} is parageometric $\lambda(\phi^{-1}) > \lambda((\phi^{-1})^{-1}) = \lambda(\phi)$. \diamond

Corollary 3. *A fully irreducible $\phi \in \text{Out}(F_r)$ is geometric if and only if the F_r -trees T_- and T_+ are both geometric.*

Proof. If ϕ is not geometric then ϕ^{-1} is also not geometric, but if T_-, T_+ were both geometric trees then it would follow by definition that ϕ and ϕ^{-1} are both parageometric, contradicting Corollary 2. The other direction was noted above. \diamond

Corollaries 2 and 3 have been proved independently by Guirardel [Gui04], by different means.

Theorem 1 and Corollary 2 were first presented at various seminars in the Fall of 2003, including the Topology Seminar at Princeton University. Our thanks go to Baris Coskunuzer of that seminar for a question which quickly inspired the proof of the “if” direction of Corollary 3.

Sketch of the proof of Theorem 1. After some preliminaries in Section 2, in Section 3 we recall results from [BF] that characterize when T_+ is a geometric F_r -tree: this happens if and only if some positive power of ϕ has a train track representative $g: G \rightarrow G$ such that G has a unique illegal turn, g has a unique periodic Nielsen path ρ (up to reversal), and $\text{Length}(\rho) = 2 \text{Length}(G)$. When ρ exists, it is necessarily a *fixed* Nielsen path, meaning that ρ is fixed up to homotopy rel endpoints by the action of g . Also, ρ decomposes at its illegal turn into $\rho = \alpha * \beta$ where α, β are legal paths of length equal to $\text{Length}(G)$. In this situation, following [BF] we construct a 2-dimensional dynamical system $k: K \rightarrow K$ representing ϕ , which we call the *wedge model*; a detailed description of the wedge model is given in Section 3.4. The 2-complex K is obtained by attaching to G a wedge W , a triangle with one side attached along α and the other side attached along β . The unattached side of W is vertical, and each vertical segment of W has endpoints on a corresponding pair of points, one in α and one in β . The effect of $g: G \rightarrow G$ is to fold ρ by some amount, and this extends to a homotopy equivalence $k: K \rightarrow K$ whose effect on W is to collapse some vertical segments of W . The vertical segments of W form leaf segments of a measured foliation \mathcal{F}^s on K called the *stable foliation* of the wedge model $k: K \rightarrow K$. The graph G is transverse to the stable foliation, and the restriction to G of the transverse measure on the stable foliation equals the Lebesgue measure along the train track G . Following [BF], in this situation we show that the attracting tree T_+ is the dual tree of the measured foliation \mathcal{F}^s . Combining this construction with results of [BH92],

we show that ϕ is parageometric if and only if the two endpoints of ρ are distinct. Moreover, in this case there exists an edge of G which is covered exactly once by the Nielsen path ρ , and so is a free edge of the 2-complex K . This simple observation, proved in Fact 8, plays a key role in the proof of Theorem 1.

In Section 4 we carry out a detailed study of \mathcal{F}^s , the stable foliation of the wedge model. We shall show that leaves of \mathcal{F}^s are trees, but they turn out to be trees of a rather thorny variety: these leaves have lots of valence 1 vertices, occurring in the interior of edges of G that are free edges of the 2-complex K . We are particularly interested in the collection of bi-infinite lines contained in leaves of \mathcal{F}^s , which we denote $\mathcal{H}(\mathcal{F}^s)$, the *hull* of \mathcal{F}^s . We shall use properties of $k: K \rightarrow K$ to essentially identify $\mathcal{H}(\mathcal{F}^s)$ with the expanding lamination of ϕ^{-1} .

Theorem 1 is proved in Section 5 by studying the asymptotic compression rate of k on lines of $\mathcal{H}(\mathcal{F}^s)$, defined to be the exponential growth rate in n of the size of a subarc of a line of $\mathcal{H}(\mathcal{F}^s)$ that k^n collapses to a point. We compute this rate in several different ways. On the one hand, since $\mathcal{H}(\mathcal{F}^s)$ is identified with the expanding lamination of ϕ^{-1} , the asymptotic compression rate of k on lines of $\mathcal{H}(\mathcal{F}^s)$ is equal to $\lambda(\phi^{-1})$. On the other hand, the abundance of valence 1 vertices in leaves of \mathcal{F}^s shows that the asymptotic compression rate of k on lines in $\mathcal{H}(\mathcal{F}^s)$ is strictly less than the exponential growth rate in n for the size of a subtree of a leaf of \mathcal{F}^s that k^n collapses to a point. The latter rate is simply the Perron-Frobenius eigenvalue of the transition matrix for $g: G \rightarrow G$, which equals $\lambda(\phi)$. This is the culminating argument of the proof that $\lambda(\phi^{-1}) < \lambda(\phi)$.

When does $\lambda(\phi) = \lambda(\phi^{-1})$? Consider a fully irreducible $\phi \in \text{Out}(F_r)$. In the wake of our results one might wonder whether $\lambda(\phi) = \lambda(\phi^{-1})$ implies that ϕ is geometric. Here is an easy construction of counterexamples: fully irreducible outer automorphisms ϕ which are not geometric and yet which satisfy $\lambda(\phi) = \lambda(\phi^{-1})$; by Theorem 1, neither is such a ϕ parageometric.

In any group G , if $g, g' \in G$ have order 2 then gg' and $(gg')^{-1}$ are conjugate. It follows that if $\Psi, \Psi' \in \text{Out}(F_r)$ have order 2, and if $\phi = \Psi\Psi'$ is fully irreducible, then $\lambda(\phi) = \lambda(\phi^{-1})$.

For a concrete example, let $\Psi \in \text{Out}(F_3)$ be the order two element represented by the automorphism $a \mapsto b, b \mapsto a, c \mapsto c$. Let $\Psi' = \Theta\Psi\Theta^{-1}$ be a conjugate of Ψ . If the conjugating element Θ is picked randomly then one might expect that $\phi = \Psi\Psi'$ is fully irreducible and has a train track representative with no periodic Nielsen paths, and so ϕ is nongeometric. Taking Θ to be the fourth power of the outer automorphism $a \mapsto b, b \mapsto c, c \mapsto \bar{b}a$ considered earlier, and applying the train track algorithm of [BH92], one obtains the following train track map $g: G \rightarrow G$ representing ϕ . The graph G has two vertices r, q , four edges B, C, D, E , with B

from q to r , C from r to q , D from r to q , and E from q to q , and g is defined by $g(B) = C\bar{E}\bar{C}DE$, $g(C) = \bar{C}\bar{B}\bar{E}\bar{D}$, $g(D) = B$, and $g(E) = CB$. The expansion factor is $\lambda = 3.199158087\dots$

To verify that g represents a nongeometric, fully irreducible outer automorphism it is sufficient to check three things. First, the transition matrix of g is positive. Second, at each of the two vertices $v = r, q$ of G , the graph of turns taken at v is connected; this is the graph with one vertex for each oriented edge with initial vertex v , and one edge for each pair of oriented edges $E_1 \neq E_2$ with initial vertex v such that the g image of some oriented edge of G contains the subpath \bar{E}_1E_2 . Third, g has no periodic Nielsen paths, which can be checked by the following expedient. Factor g into Stallings folds, $G = G_0 \rightarrow G_1 \rightarrow \dots \rightarrow G_n = G$; we did this with $n = 8$. Let \mathcal{E}_0 be the set of length 2 edge paths in G_0 with an illegal turn; there are two such paths up to reversal, $E\bar{D}$ and $\bar{E}B$. For negative integers i define \mathcal{E}_i inductively to be a set of edge paths in G_i (with indices taken modulo n) each with one illegal turn, as follows: for each $\gamma \in \mathcal{E}_i$, take all paths in G_{i-1} with exactly one illegal turn whose straightened image in G_i is γ , and put each such path in \mathcal{E}_{i-1} . Carrying this process out, we computed that the set \mathcal{E}_{-12} is empty. This shows that g has no Nielsen paths, and so it represents a fully irreducible, nongeometric outer automorphism.

Just as a check, we also inverted the sequence of Stallings folds and applied the train track algorithm to verify that the expansion factor of the inverse is also equal to $3.199158087\dots$

To get wider classes of examples, consider a fully irreducible $\phi \in \text{Out}(F_r)$ with expanding lamination denoted Λ^u . By Section 2 of [BFH97] the group $\text{Out}(F_r)$ acts on the set of expanding laminations of fully irreducible elements, and there is a homomorphism $\ell^u: \text{Stab}(\Lambda^u) \rightarrow \mathbf{R}_+$ with discrete image and finite kernel such that $\ell^u(\phi) = \lambda(\phi)$, and $\ell^u(\Psi) = \lambda(\Psi)$ as long as $\ell^u(\Psi) \geq 1$. Applying this to ϕ^{-1} with expanding lamination denoted Λ^s , we obtain a homomorphism $\ell^s: \text{Stab}(\Lambda^s) \rightarrow \mathbf{R}_+$ such that $\ell^s(\phi^{-1}) = \lambda(\phi^{-1})$ and $\ell^s(\Psi) = \lambda(\Psi)$ as long as $\ell^s(\Psi) \geq 1$. Applying Proposition 2.16 of [BFH00] it follows that $\text{Stab}(\Lambda^s) = \text{Stab}(\Lambda^u)$, a subgroup of $\text{Out}(F_r)$ that we denote V_ϕ . From the properties of the homomorphisms $\ell^u, \ell^s: V_\phi \rightarrow \mathbf{R}_+$ it follows that the infinite cyclic group $\langle \phi \rangle$ has finite index in V_ϕ , and so any two elements of V_ϕ not contained in the common kernel of ℓ^s, ℓ^u have nonzero powers that are equal. This implies that if $\lambda(\phi) = \lambda(\phi^{-1})$ then $\lambda(\Psi) = \lambda(\Psi^{-1})$ for any $\Psi \in V_\phi$. Note that V_ϕ is the *virtual centralizer* of $\langle \phi \rangle$ in $\text{Out}(F_r)$, consisting of all $\Psi \in \text{Out}(F_r)$ that commute with some positive power of ϕ ; V_ϕ is contained in the virtual centralizer because $\langle \phi \rangle$ has finite index in V_ϕ ; and if $\Psi \notin V_\phi$ then $\Psi\phi^k\Psi^{-1} \neq \phi^k$ because their attracting fixed points in $\partial\mathcal{X}_r$ are distinct, by Proposition 2.16 of [BFH00].

It might be interesting to find necessary and sufficient conditions for the condition $\lambda(\phi) = \lambda(\phi^{-1})$, for fully irreducible $\phi \in \text{Out}(F_r)$. For example, is it necessary that V_ϕ contains a fully irreducible element that is either geometric or conjugate to its

own inverse?

At the very least, it would seem that the property $\lambda(\phi) \neq \lambda(\phi^{-1})$ is generic, and the property that ϕ be fully irreducible with nongeometric fixed trees T_- , T_+ is also generic. We invite the reader to take a random word of, say, twenty or more Nielsen generators of $\text{Out}(F_3)$ and verify that the resulting outer automorphism ϕ is fully irreducible, neither geometric nor parageometric, and satisfies $\lambda(\phi) \neq \lambda(\phi^{-1})$.

2 Preliminaries

2.1 Outer automorphisms and outer space.

The definitions in this section follow several sources. For the foundations of marked graphs, **R**-trees, and outer space, including many of the facts recalled below without citation, see [CM87] and [CV86]. For concepts of irreducibility see [BH92]. A good overview is given in [Vog02].

Outer automorphisms of free groups. Fix an integer $r \geq 2$, let F_r denote the free group of rank r , let $\text{Out}(F_r) = \text{Aut}(F_r)/\text{Inn}(F_r)$ denote its outer automorphism group, and let \mathcal{C} denote its set of nontrivial conjugacy classes. Let R_r denote the rose with r -petals and identify $\pi_1(R_r) \approx F_r$, so the group $\text{Out}(F_r)$ is identified with the group of homotopy classes of self-homotopy equivalences of R_r . Given $\phi \in \text{Out}(F_r)$ let $f_\phi: R_r \rightarrow R_r$ be a representative homotopy equivalence. $\text{Out}(F_r)$ acts naturally on \mathcal{C} , and on conjugacy classes of subgroups of F_r . We say that $\phi \in \text{Out}(F_r)$ is *reducible* if there exists a nontrivial free factorization $F_r = A_1 * \cdots * A_k * B$ so that ϕ permutes the conjugacy classes of A_1, \dots, A_k . If ϕ is not reducible then it is *irreducible*. If ϕ^n is irreducible for all $n \geq 1$ then ϕ is *fully irreducible*¹. Note that ϕ is irreducible if and only if ϕ^{-1} is, and the same for complete irreducibility.

Outer space and its boundary. An F_n -tree is an **R**-tree T equipped with an action of F_n that is minimal (no proper nonempty subtree is invariant) and nonelementary (T is not a point or a line). An F_r -tree is *proper* if the action is properly discontinuous, and it is *simplicial* if T is a simplicial complex. Two F_n trees are *isometrically* (resp. *homothetically*) *conjugate* if there is an isometry (resp. homothety) between them that conjugates one action to the other. *Outer space* \mathcal{X}_r is the set of homothetic conjugacy classes of proper, simplicial F_n -trees, with topology induced by embedding $\mathcal{X}_r \rightarrow \mathcal{PR}^{\mathcal{C}}$ as follows: first embed the set of isometric conjugacy classes into $\mathbf{R}^{\mathcal{C}}$ using translation length as a class function on F_n , and then projectivize. The

¹called “irreducible with irreducible powers” or “IWIP” in the literature

image of this embedding is precompact, and its closure and boundary are denoted $\overline{\mathcal{X}}_r$ and $\partial\mathcal{X}_r = \overline{\mathcal{X}}_r - \mathcal{X}_r$.

Points of outer space can also be represented as marked graphs, as follows. A marked graph is a graph G with all vertices of valence ≥ 3 , equipped with a path metric, and with a homotopy equivalence $R_r \rightarrow G$ called the *marking*. If a base point $p \in G$ happens to be imposed, the homotopy class of a marking of G determines and is determined by an isomorphism $F_r \rightarrow \pi_1(G, p)$ up to precomposition by an inner automorphism of F_r . Two marked graphs G, G' are *isometric* (resp. *homothetic*) if there exists an isometry (resp. homothety) $G \rightarrow G'$ which, together with the markings, makes the following diagram commute up to homotopy:

$$G \xleftarrow{\quad} R_r \xrightarrow{\quad} G'$$

Passage to the universal covering space induces a bijection between the set of homothety classes of marked graphs and the set \mathcal{X}_r . The embedding $\mathcal{X}_r \rightarrow \mathcal{PR}^C$ can be understood by first associating to a marked graph G the class function on F_n that associates to an element of F_n the length of the shortest loop in G representing the free homotopy class of that element, and then projectivizing.

The length of an object in a geodesic metric space is denoted $\text{Length}(\cdot)$, with a subscript to denote the metric space when the context is not clear, for example $\text{Length}_G(\cdot)$ in the marked graph G . Also, when a marked graph G is clear from the context then we use the constant L to denote $\text{Length}(G)$.

The group $\text{Out}(F_r)$ acts on $\overline{\mathcal{X}}_r$ on the right, as follows. Let $[\cdot]$ denote homothety class. For each F_r -tree T one can precompose the action $F_r \rightarrow \text{Isom}(T)$ with an automorphism $F_r \rightarrow F_r$ representing ϕ , to get $[T]\phi$. In terms of a marked graph G , one can precompose the marking $R_r \rightarrow G$ with a homotopy equivalence $R_r \rightarrow R_r$ that represents ϕ , to obtain $[G]\phi$. This action preserves the topology, and it preserves outer space \mathcal{X}_r itself and its boundary $\partial\mathcal{X}_r$.

Source-sink dynamics. If $\phi \in \text{Out}(F_r)$ is fully irreducible then there exist $T_- \neq T_+ \in \partial\mathcal{X}_r$ such that for every $x \in \overline{\mathcal{X}}_r$,

$$\lim_{n \rightarrow -\infty} \phi^n(x) = T_- \quad \text{and} \quad \lim_{n \rightarrow +\infty} \phi^n(x) = T_+$$

where these limits take place in $\overline{\mathcal{X}}_r$. This was proved for $x \in \mathcal{X}_r$ in [BFH97] and extended to all $x \in \overline{\mathcal{X}}_r$ in [LL03]. We call T_- the *repelling tree* and T_+ the *attracting tree* of ϕ .

Geometric F_r -trees. We review measured foliations on 2-complexes and geometric trees following [LP97].

Consider a connected simplicial 2-complex K which is not a point. A *measured foliation* on K is defined by specifying measured foliations on each 2-simplex of K which fit together compatibly along 1-simplices. To be precise, a measured foliation on a 1 or 2 dimensional simplex σ in K is determined by choosing a simplicial homeomorphism $\sigma' \rightarrow \sigma$ where σ' is a rectilinear simplex in \mathbf{R}^2 , and pushing forward the vertical foliation on \mathbf{R}^2 with the transverse measure $|dx|$. A *leaf segment* in σ is the pushforward of σ' intersected with a vertical line; for example, if σ is a 1-simplex then either σ is a single leaf segment or each point of σ is a leaf segment. A measured foliation on a 2-simplex restricts to a measured foliation on each of its edges. A measured foliation on K is determined by choosing a measured foliation on each 2-simplex of K , so that for each 1-simplex e , all of the measured foliations on e obtained by restricting to e from a 2-simplex incident to e agree with each other.

We will often suppress the simplicial structure on K , so a measured foliation on a cell complex means, formally, a measured foliation on some simplicial subdivision.

Let \mathcal{F} denote a measured foliation on K , and $\mathcal{F}|_\sigma$ its restriction to each simplex σ of K . The collection of leaf segments in 1-simplices and 2-simplices define a relation on K , two points being related if they are contained in the same leaf segment. This relation generates an equivalence relation on K . The equivalence classes are called *leaves*. The leaf containing a point $x \in K$ can be built up inductively as follows: let ℓ_1 be the union of all leaf segments containing x ; for $i \geq 1$ let ℓ_{i+1} be the union of all leaf segments containing points of ℓ_i ; finally, the leaf containing x is $\cup_{i=1}^\infty \ell_i$.

Given a measured foliation \mathcal{F} on K and a path $\gamma: I \rightarrow K$, pulling back the transverse measure locally gives a measure on I , whose integral is denoted $\int_\gamma \mathcal{F}$.

Given a finite 2-complex K with measured foliation \mathcal{F} and a surjective homomorphism $h: \pi_1(K) \rightarrow F_r$, let $\tilde{K} \rightarrow K$ be the covering space corresponding to $\ker(h)$, and let $\tilde{\mathcal{F}}$ be the lifted measured foliation on \tilde{K} . Define a pseudo-metric on \tilde{K} where $d(x, y)$ is the infimum of the transverse measures of paths from x to y . Let T be the associated metric space, whose points are the equivalence classes determined by the relation $d(x, y) = 0$. Note that if x, y are in the same leaf of \tilde{K} then $d(x, y) = 0$, but the converse need not hold in general. The action of F_r on \tilde{K} induces an isometric action of F_r on the metric space T . We assume that each 1-simplex e of \tilde{K} is a geodesic, that is, if $\partial e = \{v, w\}$ then $d(v, w) = \int_e \tilde{\mathcal{F}}$. Under this assumption, Levitt and Paulin [LP97] prove that T is an F_r -tree, called the *dual F_r -tree* of the measured foliation \mathcal{F} .

An F_r -tree T is said to be *geometric* if there exists a finite 2-complex K with measured foliation \mathcal{F} , and a surjective homomorphism $\pi_1(K) \rightarrow F_r$, such that each edge of \tilde{K} is a geodesic, and such that T is isometrically conjugate to the dual F_r -tree of \mathcal{F} .

Geometric and parageometric outer automorphisms. An outer automorphism $\phi \in \text{Out}(F_r)$ is *geometric* if there exists a compact surface S , an isomorphism $F_r \approx \pi_1 S$, and a homeomorphism $h: S \rightarrow S$, such that the outer automorphism of F_r induced by h is equal to ϕ . If ϕ is fully irreducible and geometric, then its attracting and repelling F_r -trees T_+, T_- are both geometric F_r -trees.

Consider now a fully irreducible $\phi \in \text{Out}(F_r)$ with attracting F_r -tree $T_+ \in \partial \mathcal{X}_r$. We say that ϕ is *parageometric* if T_+ is a geometric F_r -tree but ϕ is not a geometric outer automorphism.

2.2 Train tracks and laminations

The definitions in this section follow [BH92] and [BFH97].

Topological representatives and Markov partitions. Given $\phi \in \text{Out}(F_r)$, a marked graph G with marking $\mu: R_r \rightarrow G$, and a homotopy equivalence $g: G \rightarrow G$, we say that g is a *topological representative* of ϕ if g takes vertices to vertices, g is an immersion on each edge, and the following diagram commutes up to homotopy:

$$\begin{array}{ccc} R_r & \xrightarrow{f_\phi} & R_r \\ \mu \downarrow & & \downarrow \mu \\ G & \xrightarrow{g} & G \end{array}$$

and so the composition $R_r \rightarrow G \xrightarrow{g} G$ represents the point $[G]\phi \in \mathcal{X}_r$. The set of edges \mathcal{E} of G forms a Markov partition for g , meaning that for any $e, e' \in \mathcal{E}$, each component of $e \cap g^{-1}(\text{int}(e'))$ is mapped by g homeomorphically onto $\text{int}(e')$. The *transition graph* $\mathcal{T}\mathcal{G}$ of g is a directed graph whose vertex set is the set \mathcal{E} , such that for each $e, e' \in \mathcal{E}$, the directed edges from e to e' are in one-to-one correspondence with the components of $e \cap g^{-1}(\text{int}(e'))$. The *transition matrix* of g is the function $M: \mathcal{E} \times \mathcal{E} \rightarrow \mathbf{Z}$, where $M(e, e')$ equals the number of directed edges in $\mathcal{T}\mathcal{G}$ from e to e' , in other words, the number of times that $g(e')$ traverses e in either direction. Note that $M^n(e, e')$ is the number of directed paths from e to e' of length n . We say that $\mathcal{T}\mathcal{G}$, M , and $g: G \rightarrow G$ are *irreducible* if there is an oriented path from any vertex of $\mathcal{T}\mathcal{G}$ to any other vertex, equivalently, for all e, e' there exists n such that $M^n(e, e') \neq 0$. If irreducibility holds then the Perron-Frobenius theorem implies that there exists a unique $\lambda \geq 1$ such that M has a positive (right) eigenvector with eigenvalue λ . If M has the stronger property that some positive power has all positive entries then $\lambda > 1$ and a positive eigenvector is unique up to positive scalar multiple.

Train tracks. A *direction* of G at a vertex v is the germ, up to reparameterization, of an immersed path with initial point v . Each direction is uniquely represented by an oriented edge e with initial point v , but we occasionally use other paths with initial point v to represent directions. A *turn* of G at v is an unordered pair $\{e, e'\}$ of directions at v ; the turn is *nondegenerate* if $e \neq e'$, otherwise the turn is *degenerate*.

An *edge path* in G will always mean a concatenation of the form $\gamma = e_0 * e_1 * \dots * e_{k-1} * e_k$, $k \geq 0$, where e_1, \dots, e_{k-1} are oriented edges and e_0, e_k are subsegments of oriented edges. Often we say “path” when “edge path” is meant; the context should make this clear. Given an edge path $\gamma: I \rightarrow G$ and $t \in \text{int}(I)$ so that $\gamma(t)$ is a vertex of G , let e, e' be the two directions of γ at this point, that is: subdivide at t to obtain a concatenation $\gamma = \alpha * \beta$, let e be the direction of α at its initial point, and let e' be the direction of β at its initial point. With this notation we say that γ *takes the turn* $\{e, e'\}$ *at the parameter value* t . If t is understood then we just say that γ *takes the turn* $\{e, e'\}$.

A topological representative $g: G \rightarrow G$ acts on the set of directions and on the set of turns of G . A nondegenerate turn is *illegal* if its image under some positive power of g is degenerate, otherwise the turn is *legal*. An edge path α is *legal* if every turn taken by α is legal, in particular every legal path is immersed. Given a path α in G , let $\alpha_\#$ denote the immersed path (or constant path) which is homotopic to α rel endpoints, so if α is immersed then $\alpha_\# = \alpha$, and if α is legal then $g^n(\alpha)_\# = g^n(\alpha)$ for all $n \geq 0$.

A topological representative $g: G \rightarrow G$ of $\phi \in \text{Out}(F_r)$ is a *train track representative* of ϕ , and g is a *train track map*, if for each edge E of G , the map $g|_E: E \rightarrow G$ is a legal path, equivalently, $g^n|_E$ is an immersion for each $n \geq 1$. A train track map g is *fully irreducible* if g^n is irreducible for all $n \geq 1$.

Proposition 4 ([BH92]). *If $\phi \in \text{Out}(F_r)$ is fully irreducible then ϕ has a fully irreducible train track representative $g: G \rightarrow G$. The transition matrix M_g has a positive power, and so there exists $\lambda(g) > 1$ and vector $v: \mathcal{E} \rightarrow \mathbf{R}$, with $\lambda(g)$ unique and v unique up to a positive scalar multiple, so that $Mv = \lambda(g)v$. We also have $\lambda(\phi) = \lambda(g)$ (see Remark 1.8 of [BH92]).*

With $g: G \rightarrow G$ as in this proposition, we may assign a path metric to G , also called the Lebesgue measure on G , so that each edge e has length $\text{Length}_G(e) = v(e)$. We may then alter g on each edge e by a homotopy rel endpoints so that $g|_e$ stretches path length by a constant factor of $\lambda(g)$; the resulting map is still a train track map. The number $\lambda(g) = \lambda(\phi)$ is called the *stretch factor* of g .

Henceforth we *always assume* without comment that if $g: G \rightarrow G$ is a train track representative of a fully irreducible $\phi \in \text{Out}(F_r)$ then g is fully irreducible and g stretches path length on G by the constant factor $\lambda(\phi)$. It follows that for any legal path γ the legal path $g(\gamma)_\# = g(\gamma)$ has length equal to $\lambda(\phi) \text{Length}(\gamma)$.

The geodesic lamination of a free group. In this heading and the next we review the results from [BFH97] concerning the construction and properties of the *expanding* or *unstable* lamination $\Lambda^u(\phi)$ of a fully irreducible $\phi \in \text{Out}(F_r)$. We use here a slightly different point of view than in [BFH97], presenting laminations as Hausdorff objects rather than non-Hausdorff.

Consider a marked graph G with universal cover T . The *geodesic lamination* of T , denoted Λ_T , is the set of pairs (ℓ, x) where $\ell \subset T$ is a bi-infinite, unoriented line and $x \in \ell$, equipped with the compact open topology where a neighborhood U_ϵ of (ℓ, x) is the set of all (ℓ', x') such that $d(x, x') \leq \epsilon$ and $\ell \cap \ell'$ contains a $1/\epsilon$ neighborhood of x in ℓ and a $1/\epsilon$ neighborhood of x' in ℓ' . The *projection map* $\Lambda_T \rightarrow T$ is the map $(\ell, x) \rightarrow x$. A *leaf* of Λ_T corresponding to a bi-infinite line ℓ is the set of all (ℓ, x) such that $x \in \ell$; we shall often confuse a bi-infinite line in T with its corresponding leaf. A point of Λ_T can also be described as a geodesic embedding $\gamma: \mathbf{R} \rightarrow T$ modulo precomposition by the involution $x \leftrightarrow -x$, where the point of Λ_T corresponding to γ is the pair $(\gamma(\mathbf{R}), \gamma(0))$. The action of F_r on T induces a properly discontinuous, cocompact action on Λ_T whose quotient space, a compact lamination denoted Λ_G , is the *geodesic lamination* of G . The projection map $\Lambda_T \rightarrow T$ descends to a projection map $\Lambda_G \rightarrow G$. An element of Λ_G can also be described as a locally geodesic immersion $\mathbf{R} \rightarrow G$ modulo precomposition by the involution $x \leftrightarrow -x$ on the parameter domain \mathbf{R} . A *sublamination* of any lamination is a closed subset that is a union of leaves. By compactness of Λ_G , every sublamination of Λ_G is compact.

Given two marked graphs G, G' , any homotopy equivalence $g: G \rightarrow G'$ induces a homeomorphism $g_*: \Lambda_G \rightarrow \Lambda_{G'}$ well defined up to isotopy, defined as follows. First alter g by homotopy so that it takes vertices to vertices and is affine on each edge. Lift g to the universal covers $\tilde{g}: T = \tilde{G} \rightarrow \tilde{G}' = T'$. There is an automorphism $\Phi: F_r \rightarrow F_r$ such that g satisfies Φ -twisted equivariance, meaning that $f(g(x)) = g(\Phi(f)(x))$ for all $x \in T$, $f \in F_r$. Consider a leaf ℓ of Λ_T . Since \tilde{g} is a quasi-isometry, $\tilde{g}(\ell)$ is a quasi-geodesic embedding of \mathbf{R} , and so the image $\tilde{g}(\ell)$ has finite Hausdorff distance from some leaf that we shall denote $\tilde{g}_*(\ell)$. Define a function $\tilde{g}_\#: \Lambda_T \rightarrow \Lambda_{T'}$ that maps each leaf ℓ to $\tilde{g}_*(\ell)$, by postcomposing the map \tilde{g} with the closest point projection from $\tilde{g}(\ell)$ onto $\tilde{g}_*(\ell)$. The map $\tilde{g}_\#$ is continuous, Φ -twisted equivariant, and induces a bijection of leaves. The image of the map $\ell \rightarrow \tilde{g}(\ell)$ is the line $\tilde{g}_*(\ell)$ union a disjoint set of finite trees attached to the line, and the effect of the closest point projection is to collapse each of these finite trees to the point where it attaches to the line; it follows that the map $\tilde{g}_\#$ is *leafwise monotonic*, meaning that for each leaf ℓ the map $\ell \xrightarrow{\tilde{g}_\#} \tilde{g}_*(\ell)$ has the property that each point pre-image is an arc. We can now perturb $\tilde{g}_\#$ to get a homeomorphism \tilde{g}_* , still satisfying Φ -twisted equivariance, and \tilde{g}_* is well-defined up to Φ -twisted equivariant isotopy. It follows that \tilde{g}_* descends to the desired homeomorphism g_* , well-defined up to isotopy.

For any $g: G \rightarrow G'$ as above that preserves the marking (in the sense that the marking $R_r \rightarrow G$, postcomposed with g , is homotopic to the marking $R_r \rightarrow G'$), note that the map $\tilde{g}_*: \Lambda_G \rightarrow \Lambda_{G'}$ is natural in the sense that for any two homotopy equivalences $G \xrightarrow{g} G' \xrightarrow{g'} G''$ that preserve markings, the composition $\Lambda_G \xrightarrow{g_*} \Lambda_{G'} \xrightarrow{g'_*} \Lambda_{G''}$ is isotopic to $(g' \circ g)_*$. We are therefore justified in talking about “the” geodesic lamination Λ_r of F_r , as represented by Λ_G for any marked graph G . We are also justified in talking about a sublamination $\Lambda' \subset \Lambda_r$ of the geodesic lamination of F_r , represented as a sublamination $\Lambda'_G \subset \Lambda_G$ for any marked graph G , with the property that for any marking preserving homotopy equivalence $g: G \rightarrow G'$ we have $g_*(\Lambda'_G) = \Lambda'_{G'}$. We also say that Λ'_G is the *realization* of Λ' in the marked graph G .

Expanding laminations. Let $\phi \in \text{Out}(F_r)$ be fully irreducible. The *expanding* or *unstable lamination* of ϕ is a sublamination $\Lambda^u(\phi)$ of the geodesic lamination of F_r , defined as follows. Choose any train track representative $g: G \rightarrow G$. Choose an edge e and a periodic point $x \in \text{int}(e)$ of periodicity p . As $n \rightarrow \infty$, the maps $e \xrightarrow{g^{np}} G$ can be reparameterized as a nested sequence of isometric immersions of larger and larger subintervals of \mathbf{R} , each interval containing 0 and each immersion taking 0 to x . The union of these immersions is a bi-infinite geodesic in G , that is, a leaf of Λ_G . The closure of this leaf in Λ_G is defined to be the realization of $\Lambda^u = \Lambda^u(\phi)$ in the marked graph G , denoted Λ_G^u . Λ^u is well-defined, independent of the choice of e and x , and also independent of g meaning that for any other train track representative $g': G' \rightarrow G'$ and any marking preserving homotopy equivalence $h: G \rightarrow G'$ we have $h_*(\Lambda_G^u) = \Lambda_{G'}^u$.

The expanding lamination $\Lambda^u(\phi)$ is *minimal*, meaning that its only nonempty sublamination is itself, in other words, every leaf is dense; see Section 1 of [BFH97]. Also, the projection from Λ_G^u to G is surjective for any train track representative $g: G \rightarrow G$, because some power of g has positive transition matrix.

Note that the action of g_* on Λ_G restricts to an action on Λ_G^u which expands length by the exact factor of $\lambda(\phi)$, that is, $\text{Length}(g_*(\ell)) = \lambda(\phi) \text{Length}(\ell)$ for any leaf segment ℓ of Λ_G^u .

3 Geometric trees and the wedge model

Let $\phi \in \text{Out}(F_r)$ be fully irreducible with attracting F_r -tree T_+ . The primary goal of this section is to characterize when T_+ is a geometric tree, and to use this characterization as an opportunity for introducing the wedge model of ϕ , which will play such an important role in later sections. We will also obtain a characterization of parageometricity of ϕ . Both characterizations are stated in terms of train track representatives of ϕ and its positive powers. These characterizations are restatements

and reworkings of results in [BF] and in [BH92]. We provide full details of proof, in part because of the limited availability of [BF], but also because of our need to develop a complete description of the wedge model.

The reader who wants to skip quickly to the definition of the wedge model should first read Section 3.1 and then Section 3.4, skipping Sections 3.2 and 3.3, although Section 3.2 will be needed to understand geometricity of T_+ .

3.1 Nielsen paths

Consider a train track map $g: G \rightarrow G$. A *(fixed) Nielsen path* of g is a locally geodesic path $\rho: [a, b] \rightarrow G$ such that $\rho(a)$ and $\rho(b)$ are fixed points, and $g \circ \rho$ is homotopic rel endpoints to ρ . A *periodic Nielsen path* of g is a Nielsen path of some power g^n with $n \geq 1$. A (periodic) Nielsen path ρ is *indivisible* if it cannot be written as a nontrivial concatenation of (periodic) Nielsen paths. Assuming that g is irreducible, every indivisible periodic Nielsen path has a unique *legal decomposition* $\rho = \alpha * \bar{\beta}$ where $\alpha, \beta: [0, b] \rightarrow G$ are legal paths of equal length, and ρ takes an illegal turn at the concatenation point. Depending on the context we may write this decomposition in other forms, for example, $\rho = \alpha_1 * \bar{\alpha}_2$. For details on Nielsen paths see [BH92].

The following result is essentially proved in [BH92] using the theory of stable train tracks. We say that a surjective map $\alpha: A \rightarrow B$ is *generated by* a relation $\sigma \subset A \times A$ if the equivalence relation generated by σ has equivalence classes identical to the point pre-images of f .

Proposition 5. *If $\phi \in \text{Out}(F_r)$ is fully irreducible then there exists $k \geq 1$ and a train track representative $g: G \rightarrow G$ of ϕ^k such that one of the following holds:*

- (1) *g has no periodic Nielsen paths.*
- (2) *(a) g has an indivisible periodic Nielsen path ρ unique up to orientation reversal, ρ is a fixed Nielsen path, the image of ρ is all of G , the illegal turn taken by ρ is the unique illegal turn of g , and $\text{Length}(\rho) = 2\text{Length}(G)$.*
*(b) Letting the legal decomposition be $\rho = \alpha * \bar{\beta}$ with $\alpha, \beta: [0, \text{Length}(G)] \rightarrow G$, the map g is generated by the relation*

$$\alpha(t) \sim \beta(t) \quad \text{for } t \in [\text{Length}(G)/\lambda, \text{Length}(G)]$$

A train track map $g: G \rightarrow G$ satisfying (2) is said to be *Nielsen unique*.

Proof. We briefly review the theory of stable train tracks. Let $g: G \rightarrow G$ be a train track representative of ϕ , let $\rho = \alpha * \bar{\beta}$ be an irreducible Nielsen path, and suppose that the illegal turn of ρ is immediately folded by g . Now apply Stallings's method

of factoring g : let E_α, E_β be maximal initial oriented segments of $\bar{\alpha}, \bar{\beta}$, respectively, such that each of E_α, E_β is contained in a single edge of G , and the paths $g(E_\alpha)$ and $g(E_\beta)$ in G are the same; and then factor g as $g = h \circ f$ where $G \xrightarrow{f} G' \xrightarrow{h} G$, the graph G' is the quotient of G obtained by identifying E_α and E_β to a single arc, $f: G \rightarrow G'$ is the quotient map, and $h: G' \rightarrow G$ is induced by g under the quotient map f . The map $g' = f \circ h: G' \rightarrow G'$ is a train track representative of ϕ , obtained from g by *folding the irreducible Nielsen path* ρ . This fold is said to be *full* if either E_α or E_β is an entire edge. Note that a sufficient condition for fullness to fail is if the Nielsen path $\rho = \alpha * \bar{\beta}$ is *small* meaning that each of the paths $\alpha, \bar{\beta}$ is a subarc of some edge. We note that the indivisible periodic Nielsen paths of g and of g' are in one-to-one, periodic preserving correspondence: paths σ, σ' correspond in this way if $\sigma' = (f \circ \sigma)_\#$.

A train track representative $g: G \rightarrow G$ of ϕ is *unstable* if there exists a sequence of train track representatives $g = g_0, g_1, \dots, g_k$ such that each g_i is obtained from the previous g_{i-1} by folding an irreducible Nielsen path, and the fold from g_{k-1} to g_k is not full. If g is not unstable then g is *stable*. It follows that any train track representative obtained from a stable train track representative by folding an irreducible Nielsen path is also stable. Stability of g implies that one of the following statements holds:

- (1') g has no indivisible fixed Nielsen paths.
- (2a') g has an indivisible fixed Nielsen path $\rho = \bar{\alpha} * \beta$ unique up to orientation reversal, the image of ρ is all of G , the illegal turn in ρ is the unique illegal turn of g , and $\text{Length}(\rho) = 2 \text{Length}(G)$.

For these statements see Section 3 of [BH92], particularly Lemma 3.9 of [BH92], and for the proof that the image of ρ is all of G see the top of page 28 of [BH92].

We turn to the proof of items (1) and (2a). Let $g': G' \rightarrow G'$ be any train track representative of ϕ . From [BFH00] Lemma 4.2.5, g' has only finitely many indivisible periodic Nielsen paths. Pass to a power of g' so that all indivisible periodic Nielsen paths of g' are fixed. In [BH92], a procedure is described which, from g' , produces a stable train track representative $g: G \rightarrow G$, and so g satisfies items (1') and (2a'). But we need to replace the word “fixed” by the word “periodic” in these items. We do this by using some of the details of the stabilization procedure, described in Section 3 of [BH92], which produces g from g' . The output of the stabilization procedure is a sequence of train track representatives $g' = g_0, g_1, \dots, g_N = g$ so that for each $n = 1, \dots, N$ one of two possibilities holds. In one case, g_n is obtained from g_{n-1} by folding an indivisible Nielsen path of g_{n-1} ; in this case the indivisible periodic Nielsen paths of g_{n-1} and of g_n are in one-to-one, period preserving correspondence, as noted above. In the other case, g_n is obtained from g_{n-1} by an operation that eliminates a small indivisible Nielsen path, and the remaining indivisible periodic Nielsen paths

of g_{n-1} that are not eliminated are in one-to-one period preserving correspondence with the indivisible periodic Nielsen paths of g_n . Since all indivisible periodic Nielsen paths of g' are fixed, the same is true for each g_n . Since $g_N = g$ has at most one indivisible fixed Nielsen path, it has at most one indivisible periodic Nielsen path, which if it exists is fixed. This proves items (1) and (2a).

Now we prove item (2b). By inductively applying Stallings fold factorization to the map $g: G \rightarrow G$ we obtain a sequence of maps

$$G_0 \xrightarrow{f_1} G_1 \xrightarrow{f_2} \dots \xrightarrow{f_K} G_K \xrightarrow{h} G_0$$

where each f_i is an edge isometric fold that preserves marking, and h is a homothety. For $k = 1, \dots, K-1$ define $g_k: G_k \rightarrow G_k$ to be

$$g_k = f_k \circ \dots \circ f_1 \circ h \circ f_K \circ \dots \circ f_{k+1}$$

and define $g_K: G_K \rightarrow G_K$ to be $g_K = f_K \circ \dots \circ f_1 \circ h$. Starting with $\rho_0 = \rho$, for each $k = 1, \dots, K$ inductively define a path ρ_k in G_k by $\rho_k = (f_k \circ \rho_{k-1})_{\#} = (f_k \circ \dots \circ f_1 \circ \rho)_{\#}$.

We prove inductively that for each $k = 1, \dots, K$, the map g_k is a stable train track representative of ϕ with unique indivisible periodic Nielsen path ρ_k , and g_k is obtained from g_{k-1} by folding ρ_{k-1} — in other words, the turn folded by f_k is the illegal turn taken by ρ_{k-1} . Assuming this is true for $k-1$, the turn of G_{k-1} folded by the map $h \circ f_K \circ \dots \circ f_k: G_{k-1} \rightarrow G_0$ is an illegal turn for g_{k-1} , but by stability g_{k-1} has a unique illegal turn, namely the illegal turn taken by the Nielsen path ρ_{k-1} , and so g_k is obtained from g_{k-1} by folding ρ_{k-1} . This implies that g_k is a stable train track representative of ϕ . Moreover, $\rho_k = (f_k \circ \rho_{k-1})_{\#}$ is an indivisible Nielsen path for g_k , and by stability ρ_k is unique.

Since the Nielsen path $\rho_k = \alpha_k * \bar{\beta}_k$ is obtained from $\rho_{k-1} = \alpha_{k-1} * \bar{\beta}_{k-1}$ by folding initial oriented segments of $\bar{\alpha}_{k-1}$ and $\bar{\beta}_{k-1}$, we can describe the situation in the following manner. Noting that $\text{Length}(\alpha_0) = \text{Length}(\beta_0) = \frac{1}{2} \text{Length}(\rho_0) = \text{Length}(G)$, there exists a partition

$$0 < R_K \leq R_{K-1} \leq \dots \leq R_1 < R_0 = \text{Length}(G)$$

such that

$$\begin{aligned} \alpha_k &= f_k \circ \alpha_{k-1} \mid [0, R_k] \\ \beta_k &= f_k \circ \beta_{k-1} \mid [0, R_k] \end{aligned}$$

and the oriented segments $\overline{\alpha_{k-1} \mid [R_k, R_{k-1}]}$ and $\overline{\beta_{k-1} \mid [R_k, R_{k-1}]}$ are the segments of $\bar{\alpha}_{k-1}$ and $\bar{\beta}_{k-1}$ that are folded by f_k . In other words, f_k is generated by the

relation $\alpha_{k-1}(t) \sim \beta_{k-1}(t)$ for $t \in [R_k, R_{k-1}]$. It now follows by induction that the map $f_k \circ \cdots \circ f_1: G_0 \rightarrow G_k$ is generated by the relation $\alpha_0(t) = \beta_0(t)$ for $t \in [R_k, \text{Length}(G)]$: what one needs for the induction step is that for any composition of surjective maps $A \xrightarrow{\alpha} B \xrightarrow{\beta} C$, if the relation σ on A generates α , if the relation τ on B generates β , and if the relation τ' on A is mapped onto the relation τ by the map $\alpha \times \alpha: A \times A \rightarrow B \times B$, then the relation $\sigma \cup \tau'$ generates $\beta \circ \alpha$. Since the maps $g = g_0$ and $f_K \circ \cdots \circ f_1$ differ by the homeomorphism h_K , it follows that g is generated by the relation that generates $f_K \circ \cdots \circ f_1$, namely the relation $\alpha(t) \sim \beta(t)$ for $t \in [R_K, \text{Length}(G)]$. Since $\rho = (g \circ \rho)_\#$, since $\text{Length}(\rho) = 2\text{Length}(G)$, and since g stretches path length by λ , a short calculation shows that $R_K = \text{Length}(G)/\lambda$. \diamond

We can now state the results that characterize geometricity of T_+ and parageometricity of ϕ . These characterizations are restatements of results from [BF] and [BH92].

Proposition 6. *Suppose that $\phi \in \text{Out}(F_r)$ is fully irreducible with attracting tree T_+ , and let $g: G \rightarrow G$ be a train track representative of a positive power of ϕ satisfying case (1) or (2) of Proposition 5.*

- (1) *If g satisfies case (1) of Proposition 5 — if g has no periodic Nielsen paths — then the tree T_+ is nongeometric.*
- (2) *If g satisfies case (2) of Proposition 5 — if g is Nielsen unique — then the tree T_+ is geometric.*

When g is Nielsen unique, letting $\rho: [0, 2\text{Length}(G)] \rightarrow G$ be the Nielsen path, we have:

- (3) *ϕ is geometric if and only if ρ is a closed path. In this case ρ traverses every edge of G exactly twice.*
- (4) *ϕ is parageometric if and only if ρ is not a closed path. In this case ρ traverses some edge of G exactly once, and ρ traverses some other edge of G at least thrice.*

The proofs of items (1) and (2) will be carried out in the remainder of Section 3. First, in Section 3.2, we give a direct limit construction which produces the tree T_+ out of the train track map g ; this leads quickly to a proof of item (1) in Section 3.3. Then, in Section 3.4, we describe the wedge model which applies to the Nielsen unique case; this leads to the proof of item (2) in Section 3.5.

We start with:

Proof of items (3) and (4), assuming (1) and (2). Let $g: G \rightarrow G$ be a Nielsen unique train track representative of a positive power of ϕ . Item (3) was proved in [BH92]. Combined with item (2), it immediately follows that ϕ is parageometric if and only if ρ is not closed.

Suppose that ϕ is parageometric. If ρ traverses every edge of G at least twice then using $\text{Length}(\rho) = 2 \text{Length}(G)$ it follows that ρ traverses every edge exactly twice, but that implies that ρ is closed, a contradiction. It follows that ρ traverses some edge E at most once. On the other hand, by Proposition 5 item (2) the path ρ traverses each edge of G at least once, and so ρ traverses E exactly once. Using again that $\text{Length}(\rho) = 2 \text{Length}(G)$ it follows that ρ traverses some other edge at least thrice. This proves item (4). \diamond

3.2 Direct limits

Let $\phi \in \text{Out}(F_r)$ be fully irreducible with attracting tree T_+ , and let $g: G \rightarrow G$ be any train track representative of ϕ . We describe here a method for constructing T_+ from g . From a more well known point of view, T_+ is the Gromov–Hausdorff limit of the universal covering trees T_0, T_1, T_2, \dots of the sequence of marked graphs $G = G_0, G_1, G_2, \dots$ given inductively by $[G_i] = [G_{i-1}]\phi$. Here we shall recast this point of view, making use instead of direct limits. In this section we give a preliminary description of the direct limit of T_i in the category of semimetric spaces and distance nonincreasing maps; this direct limit is denoted $T^\#$. Then, in Section 3.3 we study the case that g satisfies item (1) of Proposition 5, and in Sections 3.4 and 3.5 we study the case that g satisfies item (2); in both cases we will verify that the metric space associated to the semimetric space $T^\#$ is the direct limit of T_i in the category of metric spaces and distance nonincreasing maps, and we will use this information to identify the metric space direct limit with T_+ .

Consider the sequence in \mathcal{X}_r defined by $[G_0] = [G]$, $[G_i] = [G_{i-1}]\phi$. The underlying unmarked and unmetrized graphs $G = G_0, G_1, G_2, \dots$ are all identical, and the map $g: G \rightarrow G$ is rewritten as $g_i: G_i \rightarrow G_{i+1}$. We then define a marking $R_r \rightarrow G_i$, given inductively by postcomposing the marking $R_r \rightarrow G_{i-1}$ with the map g_{i-1} . It follows that each map $g_i: G_i \rightarrow G_{i+1}$ respects markings (up to homotopy). We also define a metric on G_i , given inductively as the unique metric such that g_{i-1} maps each edge of G_{i-1} locally isometrically to G_i . It follows that $\text{Length}(G_{i+1}) = \text{Length}(G_i)/\lambda(\phi)$. There is a homothety $h_i: G \rightarrow G_i$ that compresses Lebesgue measure by a factor of $\lambda(\phi)^{-i}$ such that the following composition equals g^i :

$$G = G_0 \xrightarrow{g_0} G_1 \xrightarrow{g_1} \dots \xrightarrow{g_{i-1}} G_i \xrightarrow{h_i^{-1}} G \quad (3.1)$$

Note that since the maps g_0, \dots, g_{i-1} all respect marking, the map $g_{i-1} \circ \dots \circ g_0 =$

$h_i \circ g^i: G \rightarrow G_i$ also respects marking, and so the homothety $h_i^{-1}: G_i \rightarrow G$ and the map $g^i: G \rightarrow G$ change marking by exactly the same outer automorphism, namely ϕ^i .

Let T_i be the universal cover of G_i , so we may regard T_i as an F_r -tree. Choose lifts of the maps g_i to obtain a sequence of surjective equivariant maps

$$T_0 \xrightarrow{\tilde{g}_0} T_1 \xrightarrow{\tilde{g}_1} T_2 \xrightarrow{\tilde{g}_2} \dots \quad (3.2)$$

such that \tilde{g}_i maps each edge of T_i isometrically onto an arc of T_{i+1} . Each map $\tilde{g}_i: T_i \rightarrow T_{i+1}$ is therefore distance nonincreasing. Since $[T_i] = [T_{i-1}]\phi$ as points of \mathcal{X}_r , and since T_+ is the attracting point in $\overline{\mathcal{X}}_r$ of each forward orbit of ϕ on \mathcal{X}_r [BFH97], it follows that $\lim_{i \rightarrow \infty} T_i = T_+$ in $\overline{\mathcal{X}}_r$.

Let $T^\#$ denote the *direct limit* of the sequence (3.2). Set theoretically, this is the set of equivalence classes of the disjoint union of T_0, T_1, T_2, \dots , where $x_i \in T_i$ is equivalent to $x_j \in T_j$ if there exists $k \geq i, j$ such that x_i, x_j have the same image in T_k , that is, $\tilde{g}_{k-1} \circ \dots \circ \tilde{g}_i(x_i) = \tilde{g}_{k-1} \circ \dots \circ \tilde{g}_j(x_j)$. Let $[x_i] \in T^\#$ denote the equivalence class of $x_i \in T_i$. By surjectivity of the maps \tilde{g}_i , for each i each equivalence class has the form $[x_i]$ for (at least) some $x_i \in T_i$. Define a semimetric on $T^\#$ by the formula $d_{T^\#}([x_i], [y_i]) = \lim_{i \rightarrow \infty} d_{T_i}(x_i, y_i)$ where by induction we define $x_i = \tilde{g}_{i-1}(x_{i-1})$ and similarly for y_i ; the formula is clearly well-defined independent of the choice of a representative in T_0 of a given point in $T^\#$, and the limit exists because $d_{T_i}(x_i, y_i)$ is a nonincreasing sequence of positive numbers. The actions of F_r on the trees T_i induce an isometric action of F_r on $T^\#$. The map $x_i \rightarrow [x_i]$ is an F_r -equivariant surjective function $q_i^\#: T_i \rightarrow T^\#$. This argument shows that $T^\#$ is in fact the direct limit of the sequence $T_0 \xrightarrow{\tilde{g}_0} T_1 \xrightarrow{\tilde{g}_1} \dots$, in the category of semimetric spaces with an isometric F_r -action and distance nonincreasing maps which are F_r -equivariant.

In order to understand $T^\#$ more precisely we need a result from [BFH00]. Fix $i \geq 0$ and an immersed arc or circle γ_i in G_i , and inductively define $\gamma_{j+1} = (g_j \circ \gamma_j)^\#$ for $j \geq i$. As long as γ_j is a nondegenerate path, the number of illegal turns in γ_j is nondecreasing as a function of j . It follows that *either* γ_j is eventually degenerate *or* the number of illegal turns in γ_j eventually stabilizes.

Lemma 7 (Lemma 4.2.6 of [BFH00]). *With the notation as above, assume that γ_j is not eventually degenerate, and choose J so that for $j \geq J$ the number of illegal turns in γ_j is constant. Then for each $j \geq J$ the immersion γ_j is a legal concatenation of legal paths and Nielsen paths.* \diamond

By a “legal concatenation” we mean that the turn at each concatenation point is a legal turn.

3.3 No Nielsen path, strong convergence, and the proof of (1)

We are now in a position to prove item (1) of Proposition 6. Let $\phi \in \text{Out}(F_r)$ be fully irreducible and suppose that $g: G \rightarrow G$ is a train track representative of a positive power of ϕ such that g has no periodic Nielsen path.

Using Lemma 7 together with the nonexistence of Nielsen paths it follows that, given i and a sequence of immersed arcs or circles γ_j in G_j defined for all $j \geq i$ such that $\gamma_{j+1} = (g_j \circ \gamma_j)_\#$, there exists J such that *either* γ_j is legal for all $j \geq J$, *or* γ_j is degenerate for all $j \geq J$. This has several consequences.

First we show that the semimetric space $T^\#$ described in Section 3.2 is actually a metric space. Consider two distinct points of $T^\#$, represented by a pair of points in T_i that are connected by a geodesic γ_i . Inductively define $\gamma_{j+1} = (g_j \circ \gamma_j)_\#$ for $j \geq i$. Since the endpoints of γ_i map to distinct points in $T^\#$, the path γ_j is not eventually degenerate, and so it is eventually legal. It follows that $\text{Length}(\gamma_j)$ is a positive constant for $j \geq J$, and this constant equals the semimetric distance in $T^\#$ between the two given points of $T^\#$. In other words, the distance between an arbitrary pair of distinct points in $T^\#$ is positive, so $T^\#$ is a metric space.

Next, the metric space $T^\#$ is an \mathbf{R} -tree, because the defining conditions for an \mathbf{R} -tree metric are closed conditions [CM87]. Moreover, $T^\#$ is a minimal \mathbf{R} -tree. To see why, it suffices to prove that each point of $T^\#$ lies on a bi-infinite geodesic. Consider a point of $T^\#$ represented by $x_i \in T_i$. The image of x_i downstairs in G_i lies on some leaf of the expanding lamination of ϕ , because the expanding lamination realized in G_i projects surjectively to G_i . Lifting this leaf we obtain a bi-infinite legal geodesic in T_i containing x_i , and the image of this geodesic in $T^\#$ is a bi-infinite geodesic containing $[x_i]$. Since $[x_i]$ is arbitrary in $T^\#$, this proves minimality of $T^\#$.

Next, by applying Lemma 7 to loops in G_i it follows that the sequence of translation distance functions for the F_r -trees T_i converges to the translation distance function on $T^\#$, from which it immediately follows that $T^\# = T_+$.

Finally, applying Lemma 7 again to the geodesic γ_i between any two points in T_i , if these two points map to distinct points in $T^\#$ then for sufficiently large j the path γ_j is legal and so embeds isometrically in $T^\# = T_+$ under $q_i^\#$. But this is precisely the definition of *strong convergence* of the sequence of F_r -trees T_i to the F_r -tree T_+ . By the main result of [LP97] it follows that the tree T_+ is not geometric.

3.4 Definition of the wedge model and its stable foliation \mathcal{F}^s

Suppose now that $g: G \rightarrow G$ is a Nielsen unique train track representative of (a positive power of) a fully irreducible $\phi \in \text{Out}(F_r)$. Let $\rho: [0, 2L] \rightarrow G$ be the unique Nielsen path of g , where $L = \text{Length}(G)$, and let $\rho = \alpha * \beta$ be the legal decomposition of ρ .

The *wedge model* is an extension of the train track map $g: G \rightarrow G$ to a homotopy equivalence $k: K \rightarrow K$, where K is the 2-complex obtained from G by attaching a disc W to G , identifying an arc on the boundary of W with the path ρ . The description of k requires imposing on W the structure of a “wedge”. From this structure we will also obtain a measured foliation on K denoted \mathcal{F}^s , called the *stable foliation* of k . The action of k preserves leaves of \mathcal{F}^s and multiplies the transverse measure by $\lambda(\phi)$. In Section 3.5 we will exhibit geometricity of T_+ by proving that T_+ is dual to \mathcal{F}^s .

Choose locally isometric parameterizations $\alpha, \beta: [0, L] \rightarrow G$. The wedge $W = \triangle ABC$ is the triangle in \mathbf{R}^2 with vertices $A = (0, +1)$, $B = (0, -1)$, $C = (L, 0)$. The attaching maps are $(x, y) \mapsto \alpha(x)$ for $(x, y) \in \overline{AC}$, and $(x, y) \mapsto \beta(x)$ for $(x, y) \in \overline{BC}$. The 2-dimensional cell complex K is obtained by attaching W to G in this manner. The 1-skeleton of K is equal to G union the base \overline{AB} of W , with identifications $A \sim \alpha(0)$ and $B \sim \beta(0)$. The *dihedral valence* in K of a 1-cell E of K is the total number of times that E is traversed by the attaching maps of the 2-cells of K ; by definition E is a free edge if its dihedral valence equals 1. Note that \overline{AB} is a free edge of K . Also, any edge E of G has dihedral valence equal to the number of times that ρ traverses G . By collapsing W from the free edge \overline{AB} we obtain a deformation retraction of K onto G , and so we may regard the 2-complex K as being marked by the homotopy equivalence $R_r \rightarrow G \hookrightarrow K$, where $R_r \rightarrow G$ is the given marking of the marked graph G . We may therefore identify $\pi_1(K) \approx \pi_1(G) \approx \pi_1(R_r) \approx F_r$.

The measured foliation \mathcal{F}^s is induced by the vertical measured foliation on W equipped with the transverse measure $|dx|$. To check compatibility along the 1-skeleton, observe that for each edge E of G , among the segments of $\overline{AC} \cup \overline{BC}$ that map onto E , the measures obtained on E by pushing forward $|dx|$ via the attaching map all agree with the usual Lebesgue measure on E .

Now we define the extension $k: K \rightarrow K$ of $g: G \rightarrow G$. Subdivide $W = P \cup W'$ where P is the subtrapezoid of W with one base $\overline{AB} = W \cap \{x = 0\}$ and with parallel base $W \cap \{x = L/\lambda(\phi)\}$, and W' is the subwedge $W' = \overline{W} - \overline{P}$. We call W' the *collapsed subwedge* of W for reasons about to become apparent. Define k to take P onto W , stretching the x -coordinate by $\lambda(\phi)$, and for $0 \leq x_0 \leq L/\lambda(\phi)$ contracting the y -coordinate of vertical segment $P \cap \{x = x_0\}$ by a factor of $c(x_0)$, where $c(x)$ is the unique affine function satisfying $c(0) = 1$, $c(L/\lambda(\phi)) = 0$. In W' , for $L/\lambda(\phi) \leq x_0 \leq L$ the vertical segment $W' \cap \{x = x_0\}$ is mapped by k to the point $g(\alpha(x_0)) = g(\beta(x_0))$; the latter equation follows from the definition of an indivisible Nielsen path. This completes the definition of the wedge model $k: K \rightarrow K$.

We note a few facts about k which will be important in what follows, and which are simple consequences of item (2) of Proposition 5:

- $k(x) = g(x)$ for each $x \in G$, so k is continuous.
- k is a homotopy equivalence of K , and the induced outer automorphism on

$\pi_1(K) \approx F_r$ is ϕ .

- For each $x \in K$, the set $k^{-1}(x)$ is described as follows:
 - If $x \in K - G$ then $k^{-1}(x)$ is a point in $K - G$.
 - If $x \in G$ then $k^{-1}(x)$ is either a point in G or a finite, connected graph whose edges are a (uniformly) finite union of vertical segments of the subwedge W' .
 - Every vertical segment of W is eventually collapsed by some power of k , *except* for the base \overline{AB} ; this follows because in the Nielsen path $\rho = \alpha * \beta$, with $\alpha, \beta: [0, L] \rightarrow G$, for each $t \in (0, L]$ there exists $i \geq 1$ such that $g^i(\alpha(t)) = g^i(\beta(t))$.

To understand the description of $k^{-1}(x)$, note that the map k is defined by collapsing to a point each of the vertical segments of the subwedge W' . The relation described in item (2b) of Proposition 5, which generates the map g , is exactly the same as the relation of endpoint pairs of vertical segments of the collapsed subwedge W' . It follows that for each $x \in G$, the set $k^{-1}(x)$ is a finite, connected union of vertical segments of W' : the vertices of this graph are the points of $g^{-1}(x)$, a uniformly finite set, and each point of G is an endpoint of a uniformly finite number of segments of W' , so $k^{-1}(x)$ is a uniformly finite connected graph. The graph $k^{-1}(x)$ can therefore be described by picking some vertical segment ℓ_0 of W' , then inductively defining ℓ_i be the union of ℓ_{i-1} with all vertical segments of W' that touch ℓ_{i-1} , and then taking the union of the ℓ_i to obtain $k^{-1}(x)$.

Let $K_0 = K$ and let K_i be the marked 2-complex similarly obtained from the marked graph G_i by attaching a wedge along the Nielsen path. The homotopy equivalence $k: K \rightarrow K$ induces a marking preserving homotopy equivalence $k_i: K_i \rightarrow K_{i+1}$ that agrees with $g_i: G_i \rightarrow G_{i+1}$. Note that the homothety $h_i: G \rightarrow G_i$ defined earlier extends to a homeomorphism also denoted $h_i: K \rightarrow K_i$ such that the composition $K = K_0 \xrightarrow{k_0} K_1 \xrightarrow{k_1} \dots \xrightarrow{k_{i-1}} K_i \xrightarrow{h_i^{-1}} K$ equals k^i . Letting $\mathcal{F}_0^s = \mathcal{F}^s$, there is a measure foliation \mathcal{F}_i^s defined inductively on K_i as the pushforward of \mathcal{F}_{i-1}^s via the map k_{i-1} . Note that under the homeomorphism $h_i: K \rightarrow K_i$, \mathcal{F}_i^s is the pushforward of \mathcal{F}^s with transverse measure multiplied by $\lambda(\phi)^{-i}$. Lifting to universal covers, we obtain a 2-complex \tilde{K}_i containing the tree T_i , and an action of F_r on \tilde{K}_i extending the action on T_i . The map $\tilde{g}_i: T_i \rightarrow T_{i+1}$ extends to a map $\tilde{k}_i: \tilde{K}_i \rightarrow \tilde{K}_{i+1}$ that is a lift of $k_i: K_i \rightarrow K_{i+1}$. There is an F_r -equivariant measured foliation $\tilde{\mathcal{F}}_i^s$ on \tilde{K}_i that is the lift of \mathcal{F}_i^s as well as the pushforward of $\tilde{\mathcal{F}}_{i-1}^s$ via \tilde{k}_{i-1} .

3.5 Proof of (2): duality of \mathcal{F}^s and T_+

Given a fully irreducible $\phi \in \text{Out}(F_n)$ with attracting tree T_+ , in this section we prove item (2) of Proposition 6: that under the assumption of Nielsen uniqueness, T_+ is a geometric F_r -tree.

Recall the notation: $g: G \rightarrow G$ is a Nielsen unique train track representative of ϕ^i for some $i > 0$, with Nielsen path ρ . From the constructions of Section 3.4, let $k: K \rightarrow K$ be a wedge model extension of g , with wedge W attached to G along ρ to obtain K . Let \mathcal{F}^s be the stable measured foliation of k on the 2-complex K , induced by the vertical measured foliation on W .

Recall also the sequence (3.1) of marked graphs and marked homotopy equivalences $G = G_0 \xrightarrow{g_0} G_1 \rightarrow \dots$, the universal cover sequence (3.2) of F_r -trees and F_r -equivariant maps $T_0 \xrightarrow{\tilde{g}_0} T_1 \rightarrow \dots$, and $T^\#$ the direct limit of this sequence in the category of F_r -semimetric spaces and F_r -equivariant distance nonincreasing maps.

Let \mathcal{L} denote the leaf space of the measured foliation $\tilde{\mathcal{F}}^s$, with F_r -equivariant semimetric $d_{\mathcal{L}}$ defined as follows. Recall the semimetric $d_{\tilde{K}}(x, y)$ on \tilde{K} , defined by integrating the transverse measure of $\tilde{\mathcal{F}}^s$ along paths connecting x to y and taking the infimum. When x, y are in the same leaf of \mathcal{F}^s , clearly $d_{\tilde{K}}(x, y) = 0$ by integrating along a leaf segment connecting x to y . The semimetric $d_{\tilde{K}}$ therefore induces a well-defined semimetric $d_{\mathcal{L}}$ on \mathcal{L} . Also, the F_r -map from \tilde{K} to the dual \mathbf{R} -tree of \mathcal{F}^s factors as the composition of the natural quotient F_r -map $\tilde{K} \rightarrow \mathcal{L}$ and a surjective F_r -map from \mathcal{L} to the dual \mathbf{R} -tree.

We shall prove geometricity of T_+ by proving that T_+ is equivariantly isometric to the dual F_r -tree of \mathcal{F}^s — the metric space associated to the semimetric $d_{\tilde{K}}$ on \tilde{K} . We proceed indirectly, working with the semimetric space $T^\#$ instead of the tree T_+ , and with the leaf space \mathcal{L} of $\tilde{\mathcal{F}}^s$ instead of the dual R_r -tree of \mathcal{F}^s .

Define an F_r -equivariant surjective map $\alpha: T^\# \rightarrow \mathcal{L}$ as follows: for $x \in T_0$ representing $[x] \in T^\#$, $\alpha[x]$ is the leaf of $\tilde{\mathcal{F}}^s$ passing through x . To see that α is well defined, note that for $x, y \in T_0$ we have $[x] = [y]$ if and only if x, y have the same image in some T_i , which occurs if and only if there exists a sequence of vertical segments of the wedge W connecting x to y each of which are eventually collapsed by \tilde{k} , but this implies that x, y are in the same leaf of $\tilde{\mathcal{F}}^s$.

We claim that:

(A) The map $\alpha: T^\# \rightarrow \mathcal{L}$ is distance preserving.

To prove this, consider $x_0, x'_0 \in T_0$ contained in respective leaves $\ell, \ell' \in \mathcal{L}$ of the measured foliation $\tilde{\mathcal{F}}^s$, and let $\xi = [x_0], \xi' = [x'_0] \in T^\#$, so $\alpha(\xi) = \ell$ and $\alpha(\xi') = \ell'$. Define inductively $x_i = \tilde{g}_{i-1}(x_{i-1})$ and similarly for x'_i . Let γ_i be the geodesic in T_i between x_i and x'_i , so the sequence $\text{Length}_{T_i}(\gamma_i)$ is nonincreasing and has limit

$d_{T^\#}(\xi, \xi')$. Choosing $\epsilon > 0$, for sufficiently large I we have $\text{Length}_{T_I}(\gamma_I) \leq d_{T^\#}(\xi, \xi') + \epsilon$. Starting with $\rho_I = \gamma_I$, inductively define a path ρ_i in K_i for $i = I - 1, \dots, 0$ as follows: at any point where ρ_{i+1} does not pull back continuously to \tilde{K}_i , one can interpolate a leaf segment of $\tilde{\mathcal{F}}_i^s$, producing a continuous path ρ_i in \tilde{K}_i connecting x_i to x'_i , such that $\int_{\rho_i} \tilde{\mathcal{F}}_i^s = \int_{\rho_{i+1}} \tilde{\mathcal{F}}_{i+1}^s$. We therefore have

$$\int_{\rho_0} \tilde{\mathcal{F}}_0^s = \int_{\rho_I} \tilde{\mathcal{F}}_I^s = \text{Length}_{T_I}(\gamma_I) < d_{T^\#}(\xi, \xi') + \epsilon$$

Letting $\epsilon \rightarrow 0$, it follows that $d_{\mathcal{L}}(\alpha(\xi), \alpha(\xi')) \leq d_{T^\#}(\xi, \xi')$. For the opposite inequality, choosing $\epsilon > 0$ let ρ_0 be a continuous path in \tilde{K} from x_0 to x'_0 so that $\int_{\rho_0} \tilde{\mathcal{F}}^s \leq d_T(\ell, \ell') + \epsilon/2$. Without increasing the integral along ρ_0 we may rewrite it as a concatenation of immersed paths in T_0 and vertical segments of wedges. Inductively define the path ρ_i in \tilde{K}_i as $\tilde{k}_{i-1} \circ \rho_{i-1}$, which has the effect of collapsing certain vertical wedge segments of ρ_{i-1} , and so $\text{Length}(\rho_i) = \text{Length}(\rho_{i-1})$. By induction, $\text{Length}(\rho_i) = \text{Length}(\rho_0)$. For sufficiently large i , say $i \geq I$, all vertical wedge segments of ρ_0 have been collapsed in ρ_i except for those which are lifts of the base \overline{AB} of the wedge W ; let γ_i be the path in T_i obtained from ρ_i by replacing each such vertical segment with the associated Nielsen path. For $i \geq I$ the number of these Nielsen paths is constant, and their length goes to zero as $i \rightarrow \infty$, and so for sufficiently large i we have

$$d_{T^\#}(\xi, \xi') \leq \text{Length}(\gamma_i) \leq \text{Length}(\rho_i) + \epsilon/2 < d_T(\ell, \ell') + \epsilon$$

This proves the claim.

We claim next that:

(B) For each $\xi, \xi' \in T^\#$, $\alpha(\xi) = \alpha(\xi')$ if and only if $d_{T^\#}(\xi, \xi') = 0$.

If $d_{T^\#}(\xi, \xi') \neq 0$ then the previous claim shows that $d_{\mathcal{L}}(\alpha(\xi), \alpha(\xi')) \neq 0$ and so $\alpha(\xi) \neq \alpha(\xi')$. To prove the converse, suppose that $d_{T^\#}(\xi, \xi') = 0$. Choosing $x_0, x'_0 \in T_0$ so that $\xi = [x_0]$, $\xi' = [x'_0]$, we must prove that the points x_0, x'_0 are contained in the same leaf of $\tilde{\mathcal{F}}^s$. Define inductively $x_i = \tilde{g}_{i-1}(x_{i-1})$ and similarly for x'_i . Let γ_i be the geodesic between x_i and x'_i in T_i . Applying Lemma 4.2.6 of [BFH00] (Lemma 7 above), for sufficiently large i , say $i \geq I$, the path γ_i is a legal concatenation of legal paths and Nielsen paths. As $i \geq I$ increases, the lengths of the legal paths in γ_i stay the same while the number of Nielsen paths is constant and their lengths go to zero. It follows that $d_{T^\#}(\xi, \xi') = 0$ only if γ_i has no legal paths for $i \geq I$, that is, γ_I is a concatenation of Nielsen paths. We may therefore connect x_I to x'_I by a path ρ_I in \tilde{K}_I entirely contained in a leaf of $\tilde{\mathcal{F}}_I^s$. Now proceeding inductively as in the earlier claim, for $i = I, I - 1, \dots, 0$ we obtain a path ρ_i in \tilde{K}_i entirely contained in a leaf of $\tilde{\mathcal{F}}_i^s$ connecting x_i to x'_i , and taking $i = 0$ it follows that $\alpha(\xi) = \alpha(\xi')$.

From the above two claims we can draw the conclusion that the semi-metric on \mathcal{L} is a metric: if $\ell, \ell' \in \mathcal{L}$ and $d_{\mathcal{L}}(\ell, \ell') = 0$ then, choosing $\xi \in \alpha^{-1}(\ell)$ and $\xi' \in \alpha^{-1}(\ell')$, Claim (A) implies that $d_{T^\#}(\xi, \xi') = 0$, and Claim (B) implies that $\ell = \alpha(\xi) = \alpha(\xi') = \ell'$. It immediately follows, follows from the definition of \mathcal{L} , that \mathcal{L} is equal to the associated metric space of the semimetric on \tilde{K} , in other words, \mathcal{L} is equal to the dual tree of \mathcal{F}^s . It also immediately follows that the map α identifies \mathcal{L} with the metric space associated to the semimetric on $T^\#$. Moreover, we know exactly which distinct pairs of points in $T^\#$ have distance zero, namely, those pairs ξ, ξ' represented by $x_0, x'_0 \in T_0$ contained in the same leaf ℓ of $\tilde{\mathcal{F}}_0^s$ so that x_0 and x'_0 are separated from each other in ℓ by the lifts of \overline{AB} in ℓ .

To complete the proof that T_+ is geometric, we must check that the translation distance function of the F_r -tree \mathcal{L} is the limit of the translation distance functions of the F_r -trees T_i , for that will identify \mathcal{L} with T_+ . Consider $c \in \mathcal{C}$ represented by an immersed closed curve γ_0 in G_0 . Inductively define $\gamma_i = (g_i \circ \gamma_{i-1})_\#$, so $\text{Length}_{G_i}(\gamma_i)$ is the translation length of c in T_i . By Lemma 7, for sufficiently large i , say $i \geq I$, γ_i is a legal concatenation of legal paths and Nielsen paths, where the number of Nielsen paths is a constant independent of i and their lengths go to zero, so $\lim(\text{Length}_{G_i}(\gamma_i))$ is equal to the total length of the portion of γ_i which is not in one of the Nielsen paths, the latter number being independent of i . Let $A_i \subset T_i$ be an axis of (any representative of) c acting on T_i . The map $A_i \rightarrow T^\# \rightarrow \mathcal{L} = T_+$ has the effect of folding each Nielsen path in A_i into a segment in $T^\#$, and mapping the rest of A_i onto the axis A_+ of c in T_+ . It follows that a fundamental domain for A_+ has length equal to the length of a fundamental domain for A_i minus the Nielsen paths, which equals $\lim(\text{Length}_{G_i}(\gamma_i))$.

This completes the proof of Proposition 6.

Local topology of the wedge model. When a fully irreducible outer automorphism $\phi \in \text{Out}(F_r)$ has a geometric attracting tree T_+ , there are some interesting connections between the behavior of ϕ and the topology of the wedge model K . For instance, when ϕ is geometric then K is a surface with one boundary component; this is proved in [BH92].

When ϕ is parageometric, we shall need the following fact, an immediate consequence of Proposition 6 (4) and the construction of the wedge model. This fact will play a crucial role in the proof of Theorem 1.

Fact 8. *If $\phi \in \text{Out}(F_r)$ is parageometric, if $g: G \rightarrow G$ is a Nielsen unique train track representative of a positive power of ϕ , and if $k: K \rightarrow K$ is the wedge model for g , then some edge of G is a free edge of K , and some edge has dihedral valence ≥ 3 . \diamond*

4 The stable foliation of the wedge model

In this section we fix a fully irreducible, parageometric outer automorphism $\phi \in \text{Out}(F_r)$, a Nielsen unique train track representative $g: G \rightarrow G$, and the wedge model $k: K \rightarrow K$ of g , and we study the stable foliation \mathcal{F}^s of k , a k -invariant foliation of the 2-complex K whose leaves are compressed by the action of k . The main result, Proposition 17, says that the set of bi-infinite lines in \mathcal{F}^s , called the *hull* of \mathcal{F}^s , can be identified with the leaves of the expanding lamination of ϕ^{-1} . The results of this section are closely related to results found in [BF]. Indeed, we believe that Proposition 17 can be proved by the Rips machine methods of [BF], but we have developed a different proof.

4.1 The stable foliation

For each $x \in K$ the leaf of \mathcal{F}^s through x is called the *stable leaf* of x , denoted \mathcal{F}_x^s . We can build \mathcal{F}_x^s up inductively: let $\ell_0 = x$, define ℓ_i inductively as the union of ℓ_{i-1} with all vertical segments of W that intersect ℓ_{i-1} , and then $\mathcal{F}_x^s = \cup_i \ell_i$. Note that \mathcal{F}_x^s is a locally finite, connected 1-complex with vertex set $\mathcal{F}_x^s \cap G$ and whose edges are vertical segments of the wedge W . An edge path in the leaf \mathcal{F}_x^s is therefore a concatenation of a consecutive sequence of vertical segments of W , and the number of such segments is called the *leafwise length* of the edge path. Any $p, q \in \mathcal{F}_x^s \cap G$ are connected by an edge path in \mathcal{F}_x^s , and the smallest leafwise length of such a path is called the *leafwise distance* between p and q . As we will see in Fact 12 below, each leaf \mathcal{F}_x^s is a tree, so the minimal edge path $[p, q]$ between vertices $p, q \in \mathcal{F}_x^s$ is unique, and the number of edges on this path is denoted $\text{Length}_K[p, q]$, called the *leafwise distance* between p and q . Notice that we do *not* measure Length_K using lengths of segments in the Euclidean triangle representation $W = \triangle ABC$. One should beware that the leafwise distance between p and q is not a continuous function, when regarded as a function on the set of ordered pairs $(p, q) \in G \times G$ such that p, q are contained a common leaf of \mathcal{F}_x^s , because as p, q vary in G the edge path $[p, q]$ could vary in such a way that it passes over the apex of W where the leafwise distance jumps discontinuously.

For each $x \in G$ the valence of the leaf \mathcal{F}_x^s at x equals the number of times that ρ passes over x except at an endpoint or the midpoint of the domain of ρ . If x is in the interior of an edge E of K , this number is just the dihedral valence of E in K . Since there are only finitely many vertices, it follows that there is a uniform upper bound for the valences of all vertices in all leaves of \mathcal{F}_x^s .

One point of confusion is the fact that, near the apex of the wedge W , the vertical segments of W get shorter and shorter, and it may seem possible that such segments could accumulate in a leaf of \mathcal{F}_x^s . This is not possible, however, because of local finiteness of the complex K : if v is the vertex of K to which the apex of W is

identified, and if U is a small neighborhood of v in K , then for every vertical segment α of W contained in U , and for every vertical segment α' of W such that α and α' share a common endpoint, the segment α' is not contained in U ; see also the end of the proof of Lemma 11. In fact one sees that the path topology on each leaf \mathcal{F}_x^s , which has as basis the path components of $\mathcal{F}_x^s \cap U$ over all open subsets $U \subset K$, is the same as the CW-topology on the simplicial complex \mathcal{F}_x^s .

From the definition of \mathcal{F}^s and the map k it is clear that k preserves the foliation \mathcal{F}^s , mapping each leaf \mathcal{F}_x^s onto the leaf $\mathcal{F}_{k(x)}^s$, inducing a bijection of leaves. The next fact (almost) justifies the terminology “stable foliation” for \mathcal{F}^s . Define the *strong stable set* of $x \in K$ to be

$$\mathcal{F}_x^{ss} = \{y \in K \mid \exists i \geq 0 \text{ s.t. } k^i(x) = k^i(y)\}$$

Fact 9 (The stable foliation). *Two points $x, y \in K$ are in the same leaf of \mathcal{F}^s if and only if there exists $i \geq 0$ such that $k^i(x) = k^i(y)$ or $k^i(x), k^i(y) \in \overline{AB}$. It follows that for each $x \in K$ we have:*

- (1) *If $\overline{AB} \not\subset \mathcal{F}_x^s$ then $\mathcal{F}_x^{ss} = \mathcal{F}_x^s$.*
- (2) *If $\overline{AB} \subset \mathcal{F}_x^s$ then the following hold.*
 - (a) *If $x \notin \text{int}(\overline{AB})$ then \mathcal{F}_x^{ss} is the component of $\mathcal{F}_x^s - \text{int}(\overline{AB})$ that contains x .*
 - (b) *If $x \in \text{int}(\overline{AB})$ then $\mathcal{F}_x^{ss} = \{x\}$.*

Proof. This follows immediately from the observation that for any vertical segment γ of the wedge W , there is a power of k collapsing γ to a point if and only if $\gamma \neq \overline{AB}$. \diamond

Fact 10 (Stable leaves are infinite). *Each leaf of \mathcal{F}^s is an infinite 1-complex.*

Proof. By Fact 9, for each $x \in G$ and each $n \geq 0$, the vertex set of the leaf of \mathcal{F}^s through x contains the set $g^{-n}(g^n(x))$, whose cardinality goes to infinity as $n \rightarrow \infty$ since the transition matrix of g is Perron–Frobenius. \diamond

4.2 The structure of stable leaves

Consider the universal covering spaces $\tilde{G} \subset \tilde{K}$ of $G \subset K$. Let $\tilde{\mathcal{F}}^s$ be the foliation of \tilde{K} obtained by lifting \mathcal{F}^s .

Given a leaf ℓ of \mathcal{F}^s , we study the structure of ℓ by considering a locally embedded finite edge path p in ℓ . Since p is an edge path, it starts and ends on \tilde{G} . Let the sequence of points of intersection of p with \tilde{G} be denoted x_0, x_1, \dots, x_J where J is the leafwise length of p . For example, one question we want to answer is whether ℓ is a tree, which is true if and only if $x_0 \neq x_J$ for all nontrivial edge paths p in ℓ ; see

Fact 12. Another question is whether ℓ is quasi-isometrically embedded in \tilde{K} , which is true if and only if the map $i \mapsto x_i$ is a quasi-isometric embedding $\mathbf{Z} \rightarrow \tilde{K}$; see Fact 14.

We establish further notation regarding p . For each $j = 1, \dots, J$, the segment of the path p from x_{j-1} to x_j is a vertical segment in a lift of the wedge, and so there exists a lift $\rho_j: [0, 2L] \rightarrow \tilde{G}$ of ρ or $\bar{\rho}$, and there exists $t_j \in [0, L)$, such that $x_{j-1} = \rho_j(L - t_j)$ and $x_j = \rho_j(L + t_j)$. Let $V_j = \rho_j(L) \in \tilde{G}$ be the point at which ρ_j makes its unique illegal turn. The geodesic $\overline{x_{j-1}x_j}$ in \tilde{G} is a path of length $2t_j$ obtained from ρ_j by truncating the initial and final segments of length $L - t_j$, and so $\overline{x_{j-1}x_j} = \alpha_j * \bar{\beta}_j$ concatenated at the illegal turn at V_j with $\alpha_j, \bar{\beta}_j$ legal paths of length t_j .

As a final remark, note in the discussion above that for each $j = 1, \dots, J - 1$, since the subpath of p from x_{j-1} to x_{j+1} in ℓ is locally embedded, the vertical wedge segments from x_{j-1} to x_j and from x_j to x_{j+1} are not inverses of each other, and so the paths $\alpha_j * \bar{\beta}_j$ and $\alpha_{j+1} * \bar{\beta}_{j+1}$ are not inverses of each other in \tilde{G} . Since G has a unique illegal turn, each point of \tilde{G} at which an illegal turn occurs has a unique illegal turn, from which it follows that $V_j \neq V_{j+1}$.

Lemma 11. *Let $J \geq 1$ and let p be a locally embedded edge path of leafwise length J in a leaf of $\tilde{\mathcal{F}}^s$. Using the notation above, for each $j = 0, \dots, J$ there exists i with $1 \leq i \leq J$ such that the geodesic $\overline{x_0x_J}$ makes an illegal turn at V_i , and the path $\overline{x_jV_i}$ is legal. Moreover, $\text{Length}(\overline{x_jV_i}) \leq L$.*

Proof. We will prove the final statement about length at the very last. The proof of the rest of the lemma is by induction on J , with $J = 1$ obvious. Supposing that the lemma is true for a certain J , we wish to prove it for $J + 1$.

Choose x_j with $0 \leq j \leq J + 1$. We may assume that $j \leq J$, because if $j = J + 1$ then we can just reverse the direction of p , reducing to the case $j = 0$. We therefore can apply the induction hypothesis to the subpath of p from x_0 to x_J , obtaining k with $1 \leq k \leq J$ such that $\overline{x_0x_J}$ has an illegal turn at V_k and such that $\overline{x_jV_k}$ is legal.

Decompose $\overline{x_0x_J} = \mu * \nu$ at the turn V_k ; no claims are made on the legality of μ or ν . The path $\overline{x_0x_{J+1}}$ is what you get by tightening $\mu * \nu * \alpha_{J+1} * \bar{\beta}_{J+1}$, which is done by cancelling a terminal segment of $\mu * \nu$ with an initial segment of $\alpha_J * \bar{\beta}_J$.

If the cancellation does not remove all of ν then V_k is still a turn in $\overline{x_0x_{J+1}}$ and we are done, by taking $V_i = V_k$. So we may assume that the terminal segment of $\mu * \nu$ which cancels is, at least, all of ν . Since α_{J+1} is legal it does not cancel with any of μ .

If not all of α_{J+1} cancels with ν then $\overline{x_0x_{J+1}}$ takes the illegal turn V_{J+1} . To show that the path $\overline{x_jV_{J+1}}$ is legal, this path is a concatenation of the legal path $\overline{x_jV_k}$ with the path $\overline{V_kV_{J+1}}$ which is legal because it is a subpath of α_{J+1} , and by construction

the turn at which these two paths are concatenated is not the illegal turn at the point V_k , and so this turn is legal. We are therefore done, taking $V_i = V_{J+1}$.

The remaining case is when $\nu = \bar{\alpha}_{J+1}$, which we show leads to a contradiction. In this case $V_k = V_{J+1}$ and so the illegal turns of $\overline{x_0 x_J} = \mu * \nu$ and $\alpha_{J+1} * \bar{\beta}_{J+1}$ at the common concatenation point are the same, because \tilde{G} has at most one illegal turn at each vertex. Also, the Nielsen paths ρ_{J+1} , ρ_k are equal up to orientation, because each is a lift of ρ or $\bar{\rho}$ and they have the same illegal turn. This implies that $(J+1) - k \geq 2$, from the final remark just before Lemma 11; this is where we use the hypothesis that p is locally embedded. If the orientations of the Nielsen paths ρ_k and ρ_{J+1} agree, then the points x_k and x_{J+1} lie on the same legal half of this Nielsen path, and so the geodesic $\overline{x_k x_{J+1}}$ is legal. But this geodesic falls under the induction hypothesis because $1 \leq J+1-k \leq J$, and so $\overline{x_k x_{J+1}}$ must have an illegal turn, a contradiction. If ρ_k , ρ_{J+1} are oriented in the opposite direction, then the path $\overline{x_k x_J}$ is legal, which also contradicts the induction hypothesis, noting that $1 \leq J-k \leq J$. The case $\nu = \bar{\alpha}_{J+1}$ therefore cannot occur, completing the induction.

Remark: Note the following consequence of what we have proved so far: any geodesic in \tilde{G} whose endpoints lie on the same leaf of $\tilde{\mathcal{F}}^s$ has at least one illegal turn.

It remains to prove that the legal path $\overline{x_j V_i}$ has length $\leq L$. Switching orientation of p if necessary, we may assume $j < i$. It suffices to show that the legal paths $\overline{x_j V_i}$ and $\overline{x_i V_i}$ have the same length, because

$$\text{Length } \overline{x_i V_i} \leq \frac{1}{2} \text{Length}(\rho_i) = \frac{1}{2} \text{Length}(\rho) = L$$

Note that $\overline{x_j x_i}$ is obtained by concatenating two legal paths in \tilde{G} , namely $\overline{x_j V_i}$ and the reverse of $\overline{x_i V_i}$, and then cancelling. The geodesic $\overline{x_j x_i}$ therefore has at most one illegal turn, but the remark above shows it has exactly one illegal turn. A similar argument shows that for each $k \geq 1$, if $\tilde{g}^k(x_j) \tilde{g}^k(x_i) = (\tilde{g}^k(\overline{x_j x_i}))_{\#}$ is nontrivial then it has exactly one illegal turn: it has at most one illegal turn because it is the straightened image of a path with exactly one illegal turn; but it has at least one illegal turn by the remark above. The proof breaks now into two cases, depending on whether x_i and x_j are identified by some power of \tilde{g} .

Case 1: $\tilde{g}^k(x_j) = \tilde{g}^k(x_i)$ for some $k \geq 1$. In this case the legal paths $\tilde{g}^k(\overline{x_j V_i})$ and $\tilde{g}^k(\overline{x_i V_i})$ are equal, so they have the same length, but g^k stretches the path length of every legal path by the same factor λ^k , and so $\text{Length } \overline{x_j V_i} = \text{Length } \overline{x_i V_i}$.

Case 2: $\tilde{g}^k(x_j) \neq \tilde{g}^k(x_i)$ for all $k \geq 1$. In this case the geodesic $\overline{\tilde{g}^k(x_j) \tilde{g}^k(x_i)}$ has exactly one illegal turn, for all $k \geq 1$. Applying Fact 9, it follows that for any sufficiently large k the points $\tilde{g}^k(x_j)$ and $\tilde{g}^k(x_i)$ are endpoints of some lift of \overline{AB} , and hence the geodesic $\overline{\tilde{g}^k(x_j) \tilde{g}^k(x_i)}$ is the corresponding lift of ρ , which we denote $\tilde{\rho}$. But $\overline{\tilde{g}^k(x_j) \tilde{g}^k(x_i)}$ is obtained by concatenating the legal path $\tilde{g}^k(\overline{x_j V_i})$ with the reverse of the legal path $\tilde{g}^k(\overline{x_i V_i})$ and then canceling, and what is left over after the cancellation is $\tilde{\rho}$ which is a concatenation of two legal paths of equal length, one an initial segment of $\tilde{g}^k(\overline{x_j V_i})$ and the other an initial segment of $\tilde{g}^k(\overline{x_i V_i})$. This proves that the legal paths $\tilde{g}^k(\overline{x_j V_i})$ and $\tilde{g}^k(\overline{x_i V_i})$ have the same length, and so $\text{Length } \overline{x_j V_i} = \text{Length } \overline{x_i V_i}$ as in Case 1. \diamond

4.3 Consequences of Lemma 11

The first consequence is:

Fact 12. *Each leaf of $\tilde{\mathcal{F}}^s$ is a tree.*

Proof. Lemma 11 shows that for each locally embedded edge path in a leaf of $\tilde{\mathcal{F}}^s$, the geodesic in \tilde{G} with the same endpoints has at least one illegal turn, and so the geodesic is nondegenerate and the endpoints are distinct. This implies that the leaf has no loops and so is a tree. \diamond

Fact 13. *Each leaf of \mathcal{F}^s is a tree.*

Proof. Let ℓ be a leaf of \mathcal{F}^s . By Fact 12, ℓ is π_1 -injective in K . If ℓ were not a tree then ℓ would contain an embedded loop. By Fact 9 the image of this loop under a sufficiently high power of the homotopy equivalence k is either a point or the segment \overline{AB} , violating π_1 -injectivity. \diamond

Consider now any F_r -equivariant proper geodesic metric on \tilde{K} , and any leaf ℓ of $\tilde{\mathcal{F}}^s$, with the simplicial structure on ℓ inducing a simplicial metric where each edge has length 1. Next we essentially prove that the injection $\ell \rightarrow \tilde{K}$ is a quasi-isometric embedding, with quasi-isometry constants independent of ℓ . What we actually need, and prove, is an interpretation of this statement that takes place entirely in \tilde{G} :

Fact 14. *For each leaf ℓ of $\tilde{\mathcal{F}}^s$, and each $x, y \in \ell \cap \tilde{G}$, letting $x = x_0, \dots, x_J = y$ be the points of $\ell \cap \tilde{G}$ from x to y in order, the map $i \mapsto x_i$ is a quasigeodesic embedding of $\{0, 1, \dots, J\}$ into \tilde{G} , with quasigeodesic constants independent of x, y .*

Proof. Let $\text{Length}(\cdot)$ denote length in \tilde{G} . Since J is arbitrary, it suffices to prove that $\text{Length}(\overline{x_0 x_J})$ is bounded above and below by an affine function of J . Since $\text{Length}(x_i x_{i+1}) \leq 2L$ we have $\text{Length}(\overline{x_0 x_J}) \leq 2L \cdot J$.

Lemma 11 gives a map f from the set $\{x_0, \dots, x_J\}$ to the set of illegal turns of the geodesic $\overline{x_0 x_J}$, such that $d(x_j, f(x_j)) \leq L$. There is an integer $\kappa \geq 1$ such that the L neighborhood of each illegal turn V_i intersects at most κ different edges. Lemma 11 also implies that each edge contains at most one of the points $\{x_0, \dots, x_J\}$, and so the L neighborhood of V_i contains at most κ of the points $\{x_0, \dots, x_J\}$. This shows that the map f is at most κ -to-one, and so its image has cardinality $\geq \frac{J+1}{\kappa}$. In other words, the geodesic $\overline{x_0 x_J}$ has at least $\frac{J+1}{\kappa}$ illegal turns. Letting $\eta > 0$ be the minimum length of an edge of G , it follows that

$$\text{Length}(\overline{x_0 x_J}) \geq \left(\frac{J+1}{\kappa} - 1\right)\eta = \frac{\eta}{\kappa}J + \eta\left(\frac{1}{\kappa} - 1\right)$$

◇

4.4 The hull of the stable foliation

The *hull* of a simplicial tree is the union of bi-infinite lines in the tree, and the hull of \mathcal{F}^s is the union of the hulls of its leaves, a closed foliated subset of K denoted $\mathcal{H}(\mathcal{F}^s)$. The hull of $\tilde{\mathcal{F}}^s$ is similarly defined, and $\mathcal{H}(\tilde{\mathcal{F}}^s)$ is equal to the total lift of $\mathcal{H}(\mathcal{F}^s)$. By Fact 14, each bi-infinite line in a leaf of $\mathcal{H}(\tilde{\mathcal{F}}^s)$ intersects \tilde{G} in a quasigeodesic embedding $\mathbf{Z} \rightarrow \tilde{G}$ and so has finite Hausdorff distance from a bi-infinite line in \tilde{G} . Our goal in this section is to identify this collection of lines with a natural extension of the lamination $\Lambda^u(\phi^{-1})$.

Fact 15. *$\mathcal{H}(\mathcal{F}^s)$ is a nonempty, closed lamination, using the compact-open topology on maps $\mathbf{R} \rightarrow K$. For each leaf ℓ of $\mathcal{H}\mathcal{F}^s$, $k(\ell)$ is also a leaf of $\mathcal{H}\mathcal{F}^s$, and this induces a bijection of the set of leaves of $\mathcal{H}\mathcal{F}^s$. The leaf ℓ' that maps to ℓ is the unique leaf contained in the set $k^{-1}(\ell)$.*

Proof. As we saw in Fact 10, each leaf has infinitely many vertices. Also, each vertex has finite valence — in fact, the valence of each vertex of each leaf is uniformly bounded, by the maximum cardinality of $\rho^{-1}(x)$ for $x \in G$. It follows that each leaf contains an infinite ray, and so there exists a sequence of longer and longer leaf segments α_i centered on a sequence of points x_i . Any limit point x of the sequence x_i lies on a bi-infinite line in some leaf. This shows simultaneously that $\mathcal{H}(\mathcal{F}^s)$ is closed and that it is nonempty.

The map k is finite-to-one on G , and it extends to K by collapsing to a point each vertical segment contained in the collapsed subwedge W' of W . It follows that each point pre-image is a connected subset of a leaf of K consisting of a union of boundedly many edges in that leaf, and since each leaf is a tree (by Fact 13), each point pre-image is a finite subtree. The intersection of ℓ with a point pre-image is therefore a finite subarc. In other words, the effect of k on ℓ is to collapse to a point

each of a pairwise disjoint collection of uniformly finite subarcs of ℓ , which implies that $k(\ell)$ is also a bi-infinite line in a leaf of \mathcal{F}^s . Thus k induces a well-defined self-map of the set of leaves of $\mathcal{H}\mathcal{F}^s$. This map is injective, because for each leaf ℓ of $\mathcal{H}\mathcal{F}^s$, the set $k^{-1}(\ell)$ maps to ℓ by collapsing a pairwise disjoint collection of uniformly bounded finite subtrees, and so $k^{-1}(\ell)$ contains a unique leaf ℓ' of $\mathcal{H}(\mathcal{F}^s)$. Since $k(\ell')$ is contained in ℓ and equals a leaf of $\mathcal{H}\mathcal{F}^s$, it follows that $k(\ell') = \ell$, showing that the map is surjective. \diamond

Now we set up notation for pushing leaves of $\mathcal{H}(\tilde{\mathcal{F}}^s)$ from \tilde{K} into the graph \tilde{G} . Consider a leaf ℓ of $\mathcal{H}(\tilde{\mathcal{F}}^s)$. Picking a base point $x_0 \in \ell \cap \tilde{G}$ and an orientation of ℓ determines an ordering of the set $\ell \cap \tilde{G}$, giving a bi-infinite sequence that we denote $\mathbf{x}(\ell) = (\dots, x_{-1}, x_0, x_1, \dots)$. This sequence has the following properties:

- (1) For each i , the geodesic $\overline{x_{i-1}x_i}$ in \tilde{G} is a subpath of a lift $\tilde{\rho} = \tilde{\alpha} * \tilde{\beta}$ of the Nielsen path ρ , consisting of equal length terminal subpaths of $\tilde{\alpha}$ and $\tilde{\beta}$.
- (2) The sequence $\mathbf{x}(\ell)$ is one-to-one.

Conversely, any sequence $\mathbf{x} = (\dots, x_{-1}, x_0, x_1, \dots)$ satisfying (1) and (2) is equal to $\mathbf{x}(\ell)$ for some leaf ℓ of $\mathcal{H}(\tilde{\mathcal{F}}^s)$ and some base point and orientation of ℓ .

Given a leaf ℓ of $\mathcal{H}(\tilde{\mathcal{F}}^s)$ and $\mathbf{x}(\ell) = (\dots, x_{-1}, x_0, x_1, \dots)$ as above, by Fact 14 the sequence $\mathbf{x}(\ell)$ is a quasi-isometric embedding $\mathbf{Z} \rightarrow \tilde{G}$ with uniform constants independent of ℓ . It follows that $\mathbf{x}(\ell)$ follow travels a unique bi-infinite line in \tilde{G} denoted $\bar{\ell}$. We thus obtain sublaminations of the geodesic laminations of \tilde{G} and of G :

$$\begin{aligned}\tilde{\Lambda}_{\mathcal{H}} &= \{\bar{\ell} \mid \ell \text{ is a leaf of } \mathcal{H}(\tilde{\mathcal{F}}^s)\} \subset \Lambda_{\tilde{G}} \\ \Lambda_{\mathcal{H}} &= \text{the projection of } \tilde{\Lambda}_{\mathcal{H}} \text{ to } \Lambda_G\end{aligned}$$

For a sequence \mathbf{x} satisfying (1) and (2) above, the image of \mathbf{x} under \tilde{g} also has the structure of a sequence satisfying (1) and (2), and \tilde{g} induces a quasi-isometry from \mathbf{x} to $\tilde{g}(\mathbf{x})$. It follows that \tilde{g} induces a quasi-isometry $\bar{\ell} \rightarrow \bar{\ell}'$ well-defined up to bounded distance. We may therefore set $\tilde{\phi}(\bar{\ell}) = \bar{\ell}'$, which induces a map $\Lambda_{\mathcal{H}} \rightarrow \Lambda_{\mathcal{H}}$ downstairs, that is, $\Lambda_{\mathcal{H}}$ is an invariant sublamination of the action of ϕ on the geodesic lamination of G .

A leaf ℓ of the geodesic lamination of F_r is *birecurrent* if its realization in any (equivalently, some) marked graph has the property that each finite subpath occurs infinitely often in both ends of ℓ . Equivalently, ℓ is contained in the set of limits points of each of its ends, where the convergence takes place in the geodesic lamination. Compactness of K implies that $\mathcal{H}(\mathcal{F}^s)$ is compact and so, as a consequence of the Hausdorff maximum principle, $\mathcal{H}(\mathcal{F}^s)$ has a nonempty minimal sublamination. Each leaf of a minimal sublamination is birecurrent. We therefore have shown:

Fact 16. $\Lambda_{\mathcal{H}}$ contains a birecurrent leaf. \diamond

Here is the main result of this section. Following the proof we will strengthen it by sketching how to identify $\Lambda_{\mathcal{H}}$ completely. Recall that $\Lambda^u(\phi)$ denotes the expanding lamination of ϕ and $\Lambda^u(\phi^{-1})$ is the expanding lamination of ϕ^{-1} . In the next proposition we consider the realizations of each of these laminations in the train track map $g: G \rightarrow G$ representing ϕ .

Proposition 17. $\Lambda_{\mathcal{H}} \supset \Lambda^u(\phi^{-1})$.

Proof. First we show that there is a uniform bound to the length of any legal subpath of any leaf of $\Lambda_{\mathcal{H}}$. To see why, fix a leaf $\bar{\ell}$ of $\Lambda_{\mathcal{H}}$, and let $\mathbf{x}(\ell) = (\dots, x_{-1}, x_0, x_1, \dots)$. For each $i \neq j$, each point of the geodesic $\overline{x_i x_j}$ has distance $\leq L$ from some x_k with $i \leq k \leq j$, and by applying Lemma 11 it follows that x_k has distance $\leq L$ from some illegal turn of $\overline{x_i x_j}$, so the illegal turns on $\overline{x_i x_j}$ are spaced no more than $4L$ apart. Any finite subsegment of $\bar{\ell}$ of length $> 4L$ is contained in some $\overline{x_i x_j}$ and so contains an illegal turn.

Recall from [BFH00] that a leaf ℓ of the geodesic lamination is *weakly attracted* to $\Lambda^u(\phi)$ if any finite subpath of a leaf of $\Lambda^u(\phi)$ is a subpath of $g^n(\ell)$ for some large n . Since the subpath can be a legal segment of arbitrarily large length, it follows that $g^n(\ell)$ contains legal paths of arbitrarily large length. But we have shown that if $\ell \in \Lambda_{\mathcal{H}}$ then $g^n(\ell) \in \Lambda_{\mathcal{H}}$ and so there is an upper bound to the length of legal subpaths. It follows that *no* leaf of $\Lambda_{\mathcal{H}}$ is weakly attracted to $\Lambda^u(\phi)$. Applying Theorem 6.0.1 of [BFH00], it follows that each birecurrent leaf of $\Lambda_{\mathcal{H}}$ is a leaf of $\Lambda^u(\phi^{-1})$, so there exists at least one leaf of $\Lambda^u(\phi^{-1})$ that is contained in $\Lambda_{\mathcal{H}}$. By minimality of $\Lambda^u(\phi^{-1})$ it follows that $\Lambda^u(\phi^{-1}) \subset \Lambda_{\mathcal{H}}$. \diamond

For completeness sake we give a description of $\Lambda_{\mathcal{H}}$, but with details of proof only sketched lightly since we do not need this description for our present purposes. We'll assume that $\bar{g}: \bar{G} \rightarrow \bar{G}$ is a train track representative of ϕ^{-1} which is either Nielsen unique or has no Nielsen path at all. The *extended expanding lamination* of ϕ^{-1} is defined to be the union of $\Lambda^u(\phi^{-1})$ with finitely many other leaves, as follows. At a periodic vertex p of \bar{g} , any periodic direction d determines a ray r_d under iteration of \bar{g} , and any two periodic directions d, d' determine a leaf $\ell = r_d \cup r_{d'}$ in the extended expanding lamination. Also, if ρ is the Nielsen path connecting two points p and q , then for any periodic directions d, d' at p, q respectively, distinct from the directions of ρ at its endpoints, $\ell = r_d \cup \rho \cup r_{d'}$ is a leaf in the extended expanding lamination. A study of the stabilization algorithm of [BH92] shows that the extended expanding lamination is well-defined, independent of the choice of \bar{g} .

We claim that $\Lambda_{\mathcal{H}}$ is the extended expanding lamination of ϕ^{-1} . This is a consequence of the proof of Theorem 6.0.1 of [BFH00] applied to a train track representative

$g: G \rightarrow G$ for ϕ . This proof is laid out with Steps 1, 2, and 3. In the end of Step 2, one considers a geodesic ℓ , for example any leaf of $\Lambda_{\mathcal{H}}$, whose realization in G is not weakly attracted to $\Lambda^u(\phi)$, meaning that the sequence $(g^i \ell)_{\#}$ does not develop longer and longer segments that are leaf segments of $\Lambda^u(\phi)$. What one shows in this situation is that there exists an immersed loop γ in \tilde{G} of uniformly bounded length so that the sequence $\gamma_i = (\bar{g}^i \gamma)_{\#}$ develops longer and longer segments that are contained in and exhaust ℓ , as realized in \tilde{G} . Applying Lemma 7, for some I the loop γ_I is a legal concatenation of legal paths and Nielsen paths, and it follows that for $i \geq I$ the loop γ_i is a legal concatenation of a uniformly bounded number of segments each of which is either a leaf segment of $\Lambda^u(\phi^{-1})$ or of one of the finitely many leaves added to make the extended expanding lamination of ϕ^{-1} . One then sees that ℓ is exhausted by such segments, implying that ℓ itself is a leaf of the extended expanding lamination.

5 Proof of Theorem 1

Recall the notation: $\phi \in \text{Out}(F_r)$ is parageometric, with $g: G \rightarrow G$ a Nielsen unique train track representative of some positive power of ϕ , with wedge model $k: K \rightarrow K$, with \mathcal{F}^s the stable foliation of k , and with $\Lambda_{\mathcal{H}}$ the sublamination of Λ_G obtained by pushing leaves of $\mathcal{H}(\mathcal{F}^s)$ into G .

To prove that $\lambda(\phi^{-1}) < \lambda(\phi)$ we will make use of an intermediate quantity, the asymptotic compression factor of k acting on leaves of $\mathcal{H}(\mathcal{F}^s)$. On the one hand we show that $\lambda(\phi)$ is a strict upper bound for this factor, by using a symbolic dynamics argument to interpret $\lambda(\phi)$ geometrically, and by exploiting Fact 8 which says that some edge of G is a free edge in K . On the other hand, we use the result from Section 4, that $\Lambda_{\mathcal{H}}$ is a sublamination of $\Lambda^u(\phi^{-1})$, to show that the asymptotic compression factor is equal to $\lambda(\phi^{-1})$ on the nose.

In this section we distinguish two measurements of length. First, we use $\text{Length}_K(\alpha)$ to denote combinatorial length of an edge path α in a leaf of \mathcal{F}^s or $\tilde{\mathcal{F}}^s$, that is, the number of vertical wedge segments in α . We also use Length_G to denote length of a path in the graph G or \tilde{G} .

By Fact 8, there is at least one edge of G that is a free edge of the 2-complex K . Let G_1 be the subgraph of G consisting of the nonfree edges, that is, the edges of dihedral valence > 1 in K . By Fact 8 it follows that G_1 is a nonempty, proper subgraph. Let $I_1 = \{x \in G \mid k^n(x) \in G_1 \text{ for all } n \geq 0\}$. Since g has constant stretch factor $\lambda(\phi)$, since the transition matrix of g has a positive power, and since G_1 is a nonempty, proper subgraph of G , it follows that I_1 is a Cantor set in G_1 .

Fact 18. *If ℓ is a leaf of $\Lambda_{\mathcal{H}}$ then $\ell \cap G \subset I_1$.*

Proof. The hypothesis means that ℓ is a bi-infinite line in a leaf of \mathcal{F}^s . This implies that $\ell \cap G \subset G_1$, because for each point $x \in G$ that is not a vertex, the valence of x

in its \mathcal{F}^s leaf equals the dihedral valence of the edge containing x . Since $k^n(\ell)$ is also a bi-infinite line in a leaf of \mathcal{F}^s , it follows that $k^n(\ell) \cap G \subset G_1$. This being true for all n , it follows that $\ell \cap G \subset I_1$. \diamond

Our next fact gives the strict inequality that we shall need in proving that $\lambda(\phi^{-1}) < \lambda(\phi)$.

Fact 19. *There is a number $\lambda' < \lambda(\phi)$ such that for each $x \in I_1$ we have*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log |g^{-n}(x) \cap I_1| \leq \log(\lambda')$$

Proof. Recall that \mathcal{TG} is the transition graph of G . Let \mathcal{TG}_1 be the subgraph obtained from \mathcal{TG} by throwing away each vertex associated to an edge of G of dihedral valence $= 1$, and any directed edge of \mathcal{TG} incident to such a vertex. Let M_1 be the transition matrix of \mathcal{TG}_1 . If x is contained in the interior of the i^{th} edge of G , then $|g^{-n}(x) \cap I_1|$ is the sum of the entries in the i^{th} column of M_1^n . In all cases $|g^{-n}(x) \cap I_1|$ is bounded above by $|M_1^n|$, the sum of all coefficients in M_1^n . It therefore suffices to show that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log |M_1^n| \leq \log(\lambda') < \log(\lambda).$$

Since we may regard M_1 as defined on the same set $\mathcal{E} \times \mathcal{E}$ as M with $M_1(e, e') \leq M(e, e')$ and with strict inequality for at least one pair $(e, e') \in \mathcal{E} \times \mathcal{E}$, and since λ is the Perron-Frobenius eigenvalue of M , this inequality follows from Perron-Frobenius theory. See for example Theorem 4.4.7 and Theorem 4.4.4 of [LM95]. \diamond

Now we relate Fact 19 to the asymptotic compression factor of k acting on leaves of \mathcal{HF}^s .

Fact 20. *For each $\epsilon > 0$ there exists an integer $N \geq 1$ such that for all $n \geq N$, if α is an arc in a leaf ℓ of \mathcal{HF}^s , and if $\text{Length}_K(\alpha)$ is sufficiently long (depending on n), then*

$$\frac{\text{Length}_K(\alpha)}{\text{Length}_K(k^n \alpha)} \leq (\lambda' + \epsilon)^n$$

Proof. Applying Fact 15, for each leaf ℓ of \mathcal{HF}^s let $\bar{k}(\ell)$ denote the unique leaf of \mathcal{HF}^s such that $k(\bar{k}(\ell)) = \ell$, so $\bar{k}(\ell) \subset k^{-1}(\ell)$. We extend the set map \bar{k} to subsets of leaves of \mathcal{HF}^s , by setting $\bar{k}(\alpha) = \bar{k}(\ell) \cap k^{-1}(\alpha)$ for each $\alpha \subset \ell$.

Let ℓ be a leaf of \mathcal{HF}^s and $\alpha \subset \ell$ an arc. Let $k^n \alpha \cap G = \{x_0, x_1, \dots, x_I\}$ in order. Note that if $0 < i < I$ then $\bar{k}^n(x_i) \subset \alpha$, so we can augment α , replacing it by $\bar{k}^n(x_0) \cup \alpha \cup \bar{k}^n(x_I)$, and increasing the length of α by an amount depending only on n , without changing $k^n \alpha$. If $\text{Length}_K(\alpha)$ were sufficiently long to start with then this

increase would change $\text{Length}_K(\alpha)$ by an arbitrarily small factor. We may therefore assume that $\alpha = \bar{k}^n(k^n\alpha)$.

It follows that α is equal to the disjoint union of the arcs $\bar{k}^n(x_0), \dots, \bar{k}^n(x_I)$ together with I additional open vertical wedge segments, each of which maps homeomorphically by k^n to the I open wedge segments that constitute the arc $k^n\alpha$. We therefore have

$$\text{Length}_K(\alpha) = \sum_{i=0}^I \text{Length}_K(\bar{k}^n(x_i)) + I$$

and dividing by $\text{Length}_K(k^n\alpha) = I$ we get

$$\frac{\text{Length}_K(\alpha)}{\text{Length}_K(k^n\alpha)} = \frac{1}{I} \sum_{i=0}^I \text{Length}_K(\bar{k}^n(x_i)) + 1$$

We also have $\text{Length}_K(\bar{k}^n(x_i)) \leq |g^{-n}(x_i) \cap I_1|$. By combining this with Fact 19 it follows that if n is sufficiently large then $\text{Length}_K(\bar{k}^n(x_i)) \leq (\lambda' + \frac{\epsilon}{2})^n$ and so

$$\frac{\text{Length}_K(\alpha)}{\text{Length}_K(k^n\alpha)} \leq \left(\lambda' + \frac{\epsilon}{2}\right)^n + \frac{\lambda' + \frac{\epsilon}{2}}{I} + 1$$

By taking n sufficiently large, this last quantity is

$$\leq (\lambda' + \epsilon)^n$$

◇

Now we prove Theorem 1. Choose a train track representative $\gamma: \Gamma \rightarrow \Gamma$ of ϕ^{-1} and let $\tilde{\gamma}: \tilde{\Gamma} \rightarrow \tilde{\Gamma}$ be a lift to the universal cover. Since $g: G \rightarrow G$ is a train track representative of ϕ , we can choose the lifts $\tilde{\gamma}: \tilde{\Gamma} \rightarrow \tilde{\Gamma}$ and $\tilde{g}: \tilde{G} \rightarrow \tilde{G}$ to represent inverse automorphisms of F_r . It follows that there is an F_r -equivariant quasi-isometry $h: \tilde{G} \rightarrow \tilde{\Gamma}$ such that h is a “quasiconjugacy” between \tilde{k} and “ $\tilde{\gamma}^{-1}$ ”, meaning that $d(\tilde{\gamma} \circ h \circ \tilde{g}(x), h(x))$ is uniformly bounded over all $x \in \tilde{G}$. Now we apply this to a particular leaf of $\Lambda^u(\phi^{-1})$ as follows. Pick a leaf ℓ' of $\Lambda^u(\phi^{-1})$ realized in Γ . For convenience we assume ℓ' is periodic under γ . By passing to a power of ϕ we may assume that ℓ' is fixed by γ . We may choose a lift $\tilde{\ell}' \subset \tilde{\Gamma}$, and we may choose the lifts \tilde{g} and $\tilde{\gamma}$, so that $\tilde{\ell}'$ is fixed by $\tilde{\gamma}$. Let $\tilde{\ell}$ denote the corresponding leaf of $\mathcal{H}\tilde{\mathcal{F}}^s$, and so $\tilde{\ell}$ is fixed by \tilde{k} . Pushing $\tilde{\ell}$ into \tilde{G} , and mapping over by h to $\tilde{\Gamma}$, the result is Hausdorff equivalent to $\tilde{\ell}'$, and composing by the closest point projection to $\tilde{\ell}'$, we obtain a map still denoted $h: \tilde{\ell} \rightarrow \tilde{\ell}'$ which is a quasi-isometry and a quasiconjugacy, that is, there are constants $\kappa \geq 1$, $\eta \geq 0$ such that

$$\frac{1}{\kappa} d_{\ell}(x, y) - \eta \leq d_{\ell'}(hx, hy) \leq \kappa d_{\ell}(x, y) + \eta, \quad \text{for all } x, y \in \tilde{\ell}$$

and such that

$$d_{\ell'}(\tilde{\gamma} \circ h \circ \tilde{k}(x), h(x)) \leq \eta \quad \text{for all } x \in \tilde{\ell}.$$

Here we use d_{ℓ} for distance in $\tilde{\ell}$ and $d_{\ell'}$ for distance in $\tilde{\ell}'$. By induction we obtain for each n a constant $\eta_n \geq 0$ such that

$$d_{\ell'}(\tilde{\gamma}^n \circ h \circ \tilde{k}^n(x), h(x)) \leq \eta_n, \quad \text{for all } x \in \tilde{\ell}'$$

We know that $\tilde{\gamma}$ expands length on $\tilde{\ell}'$ by the exact factor of $\lambda(\phi^{-1})$. On the other hand, Fact 20 tells us that \tilde{k} contracts length on $\tilde{\ell}$ by an asymptotic factor of at most λ' . Combining these, we now show that $\lambda(\phi^{-1}) \leq \lambda'$.

By Fact 20, for each sufficiently long leaf segment $[x, y]$ of $\tilde{\ell}$ we have

$$d_{\ell}(x, y) \leq (\lambda' + \epsilon)^n d_{\ell}(\tilde{k}^n(x), \tilde{k}^n(y))$$

and we also have

$$\begin{aligned} d_{\ell}(\tilde{k}^n(x), \tilde{k}^n(y)) &\leq \kappa d_{\ell'}(h \circ \tilde{k}^n(x), h \circ \tilde{k}^n(y)) + \kappa \eta \\ d_{\ell'}(h \circ \tilde{k}^n(x), h \circ \tilde{k}^n(y)) &= \lambda(\phi^{-1})^{-n} d_{\ell'}(\tilde{\gamma}^n \circ h \circ \tilde{k}^n(x), \tilde{\gamma}^n \circ h \circ \tilde{k}^n(y)) \\ d_{\ell'}(\tilde{\gamma}^n \circ h \circ \tilde{k}^n(x), \tilde{\gamma}^n \circ h \circ \tilde{k}^n(y)) &\leq d_{\ell'}(hx, hy) + 2\eta_n \\ d_{\ell'}(hx, hy) &\leq \kappa d_{\ell}(x, y) + \eta \end{aligned}$$

Combining these we get

$$\begin{aligned} d_{\ell}(x, y) &\leq (\lambda' + \epsilon)^n \lambda(\phi^{-1})^{-n} \kappa^2 d_{\ell}(x, y) + \underbrace{(\lambda' + \epsilon)^n (\eta + 2\eta_n + \kappa \eta)}_{B_n} \\ 1 &\leq (\lambda' + \epsilon)^n \lambda(\phi^{-1})^{-n} \kappa^2 + \frac{B_n}{d_{\ell}(x, y)} \end{aligned}$$

By Fact 20, we can let $d_{\ell}(x, y) \rightarrow \infty$ and we get

$$\begin{aligned} \lambda(\phi^{-1})^n &\leq (\lambda' + \epsilon)^n \kappa^2 \\ \lambda(\phi^{-1}) &\leq (\lambda' + \epsilon) \kappa^{2/n} \end{aligned}$$

and then we can let $n \rightarrow \infty$ to get

$$\lambda(\phi^{-1}) \leq \lambda' + \epsilon$$

Finally, letting $\epsilon \rightarrow 0$ we get

$$\lambda(\phi^{-1}) \leq \lambda' < \lambda(\phi)$$

This completes the proof of Theorem 1.

References

- [BF] M. Bestvina and M. Feighn, *Outer limits*, preprint.
- [BF95] ———, *Stable actions of groups on real trees*, Invent. Math. **121** (1995), no. 2, 287–321.
- [BFH97] M. Bestvina, M. Feighn, and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, Geom. Funct. Anal. **7** (1997), 215–244.
- [BFH00] ———, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms.*, Ann. of Math. **151** (2000), no. 2, 517–623.
- [BFH05] ———, *The Tits alternative for $\text{Out}(F_n)$. II. A Kolchin type theorem*, Ann. of Math. **161** (2005), no. 1, 1–59.
- [BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. **135** (1992), 1–51.
- [CL95] M. Cohen and M. Lustig, *Very small group actions on \mathbf{R} -trees and Dehn twist automorphisms*, Topology **34** (1995), no. 3, 575–617.
- [CM87] M. Culler and J. W. Morgan, *Group actions on \mathbf{R} -trees*, Proc. London Math. Soc. (3) **55** (1987), no. 3, 571–604.
- [CV86] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), 571–604.
- [FLP⁺79] A. Fathi, F. Laudenbach, V. Poenaru, et al., *Travaux de Thurston sur les surfaces*, Astérisque, vol. 66–67, Société Mathématique de France, 1979.
- [GJLL98] D. Gaboriau, A. Jaeger, G. Levitt, and M. Lustig, *An index for counting fixed points of automorphisms of free groups*, Duke Math. J. **93** (1998), no. 3, 425–452.
- [Gui04] V. Guirardel, *Core and intersection number for group actions on \mathbf{R} -trees*, preprint, arXiv:math.GR/0407206, 2004.
- [HM06a] M. Handel and L. Mosher, *Axes in outer space*, Preprint, arXiv:math.GR/0605355, 2006.
- [HM06b] ———, *The expansion factors of an outer automorphism and its inverse*, Trans. AMS (2006), To appear. Preprint, arXiv:math.GR/0410015.

- [LL03] G. Levitt and M. Lustig, *Irreducible automorphisms of F_n have north-south dynamics on compactified outer space*, J. Inst. Math. Jussieu **2** (2003), no. 1, 59–72.
- [LM95] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Chaos*, Cambridge University Press, 1995.
- [LP97] G. Levitt and F. Paulin, *Geometric group actions on trees*, Amer. J. Math. **119** (1997), no. 1, 83–102.
- [McC85] J. McCarthy, *A “Tits-alternative” for subgroups of surface mapping class groups*, Trans. AMS **291** (1985), no. 2, 582–612.
- [Vog02] K. Vogtmann, *Automorphisms of free groups and outer space*, Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000), vol. 94, 2002, pp. 1–31.

MICHAEL HANDEL:

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE

LEHMAN COLLEGE - CUNY

250 BEDFORD PARK BOULEVARD W

BRONX, NY 10468

MICHAEL.HANDEL@LEHMAN.CUNY.EDU

LEE MOSHER:

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE

RUTGERS UNIVERSITY AT NEWARK

NEWARK, NJ 07102

MOSHER@ANDROMEDA.RUTGERS.EDU